Electronic Specifier design

October 2014 | Volume 4, Issue 10

www.electronicspecifier.com

Getting to 5G

Significant challenges for seamless surfing

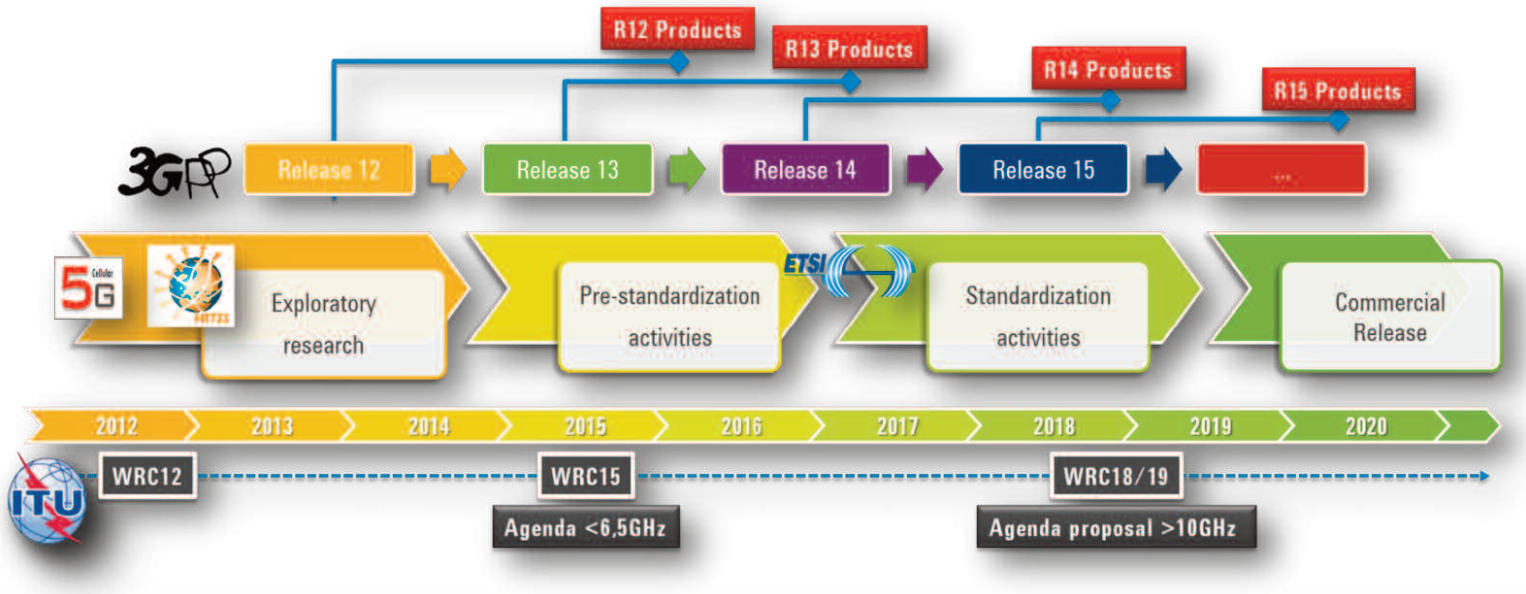

Energy

How the tides could be turning for renewable energy

Data Transfer

Will the latest version of USB provide a single connector for all devices?

Wireless

The challenges involved with delivering seamless surfing on 5G networks

Reader Offer!

Your chance to win a Microstick development platform

Farnell element14

Volume pricing
Economy with scale

We now offer larger volumes and extended price breaks on over 92,000 products

START HERE ▷

Research > Design > Production

uk.farnell.com

Electronic Specifier

design

Contents

06 News
The latest developments

10 Markets & Trends
Swinging in favour of SoC

12 Unseen Energy
Are we ready for untethered charging?

14 Harnessing the Tides
Specialised electronics fuel a tidal energy project

20 Charging towards fewer cables
The growing deployment of wireless charging

24 Batteries are included
Why battery power is a key element in creating the IoT

28 Mixing it up
How hybrid relays can simplify compliance with the latest legislation

32 Challenges of Wireless Charging
Keeping handsets topped up on the go

35 C what you get
Specifications for the new Type C connector for USB 3.1

38 Simplifying Standards
Evaluating SERDES in FPGAs

43 Getting to 5G
Significant challenges for seamless surfing

46 Getting the best out of DAS
Dealing with tomorrow's network capacity variations

49 Small Changes; Big Differences
Tackling complex designs with new functionality

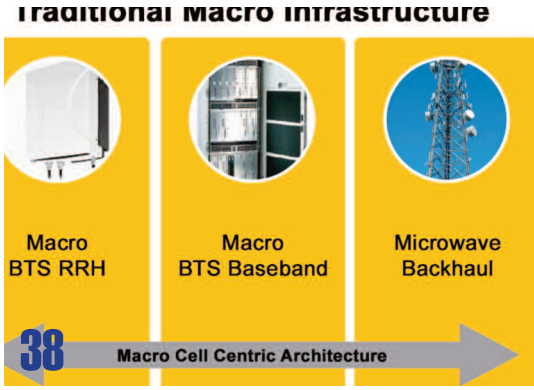

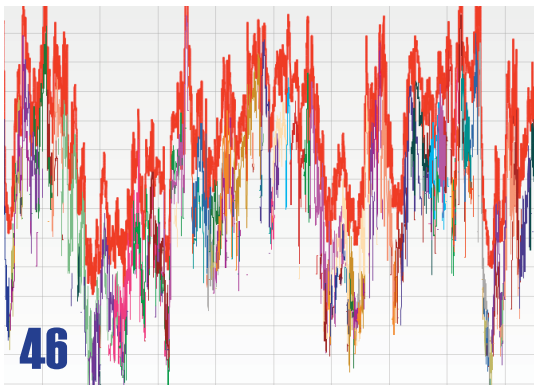

Editor:
Philip Ling
phil.ling@electronicspecifier.com

Designer:
Stuart Pritchard
stuart@origination-studio.co.uk

Ad sales:
Ben Price
ben.price@electronicspecifier.com

Publishing Director
Steve Regnier
steve.regnier@electronicspecifier.com

Head Office:
ElectronicSpecifier Ltd
Comice Place, Woodfalls Farm
Gravelly Ways, Laddingford
Kent. ME18 6DA
Tel: 01622 871944
www.electronicspecifier.com

Copyright 2013 Electronic Specifier. Contents of Electronic Specifier, its publication, websites and newsletters are the property of the publisher. The publisher and the sponsors of this magazine are not responsible for the results of any actions or omissions taken on the basis of information inthis publication. In particular, no liability can be accepted in result ofany claim based on or in relation to material provided for inclusion. Electronic Specifier is a controlled circulation journal.

Power is the new Signal

The concept of having easy access to power in public places must fill most 'power-users' with joy; we've become so dependent on our portable appliances that we can regularly suffer from 'battery anxiety' if the power bar falls below half way and we're far from home/office/airport/car…

Wireless charging would appear to be a unique situation; where demand outstrips supply. Normally in such a situation providers flood the market with products/services anxious to tap in to the waiting market, so why isn't that happening with wireless charging? The answer is probably more complex than the industry failing to converge on a single standard; the extra cost, space, weight and so on will all contribute in some way towards supplier inertia, even if it's what users want.

But, as we report in this issue, there are a number of wire-free power solutions vying for market acceptance, which promise truly untethered charging. It relies on generating envelopes of power around our devices (see page 12 for more).

I expect we all remember hanging out of a window or standing in a rainy doorway in order to get a better mobile signal, with wire-free charging seemingly just about to hit the market, could we be heading for the return of 'hunt the signal'? Wire-free charging promises to overcome the drawbacks of wireless charging, but it has its own limitations; the number of devices a single transmitter can service at any one time. This means that, in public places, we could find ourselves last in the queue for a share of the 'free' power.

The widest selection of the newest products
Over 4 million products from over 500 manufacturers.
mouser.com
Authorised Distributor
Visit mouser.com

Agilent's Electronic Measurement Group, including its 9,500 employees and 12,000 products, is now **Keysight Technologies**.

Learn more at **www.keysight.com**

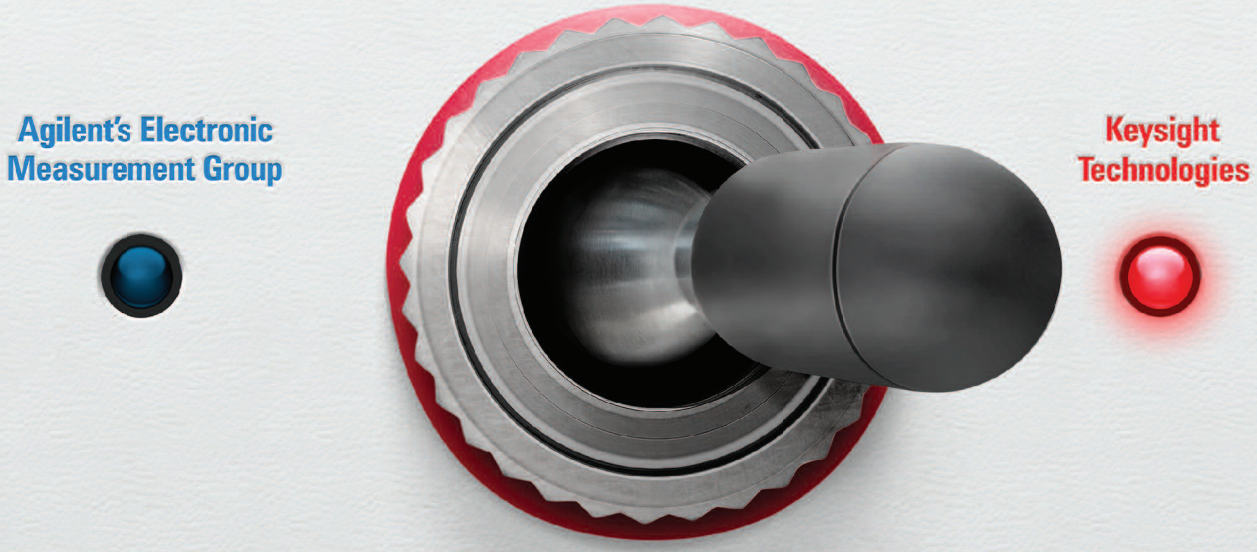

KEYSIGHT
TECHNOLOGIES

Reaching out

ARM has announced the latest version of its Cortex-M family; the -M7, offering up to twice the performance of the Cortex-M4 core. It features a six-stage superscaler pipeline and supports 64-bit transfer and optional instruction and data caches. It also offers licensees 'extensive implementation configurability' to cover a wide range of cost and performance points. ARM has also included an optional safety package, allowing it to target safety-related applications including automotive, industrial , transport and medical.

Initial licensees at launch include Freescale, Atmel and STMicro. Freescale, which was the lead partner for the Cortex-M0+ and first to market with a product, has already announced its intention to extend the Kinetis MCU family using the core. However, STMicro has claimed the title of first to market with the -M7, announcing the STM32 F7 Series which is sampling now. The family will be manufactured on ST's 90nm embedded non-volatile memory CMOS process. IAR and Segger have also announced that their tools will support the Cortex-M7.

R you IN?

In support of its own technology targeting equipment incorporating industrial Internet communication, Renesas is to establish the R-IN Consortium. It will provide global support for developers of devices such as manufacturing equipment, security cameras and robots, or other devices using Renesas' multi-protocol chips that feature the R-IN engine.

Renesas says it is now seeking corporate partners willing to supply software, operating systems, development environments and system integration services, with the intention of commencing commercial activities by April 2015.

Power-up

Supporting applications of up to 10W, Vishay has introduced a WPA-compliant wireless charging receive coil that is more than 70% efficient. Rated at 2A, the powdered-iron based coil — which is also RoHS compliant — can deliver 10W at 5V. For more on wireless charging, see the Energy section in this issue starting on page 12.

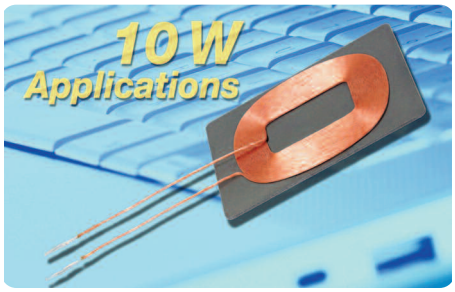

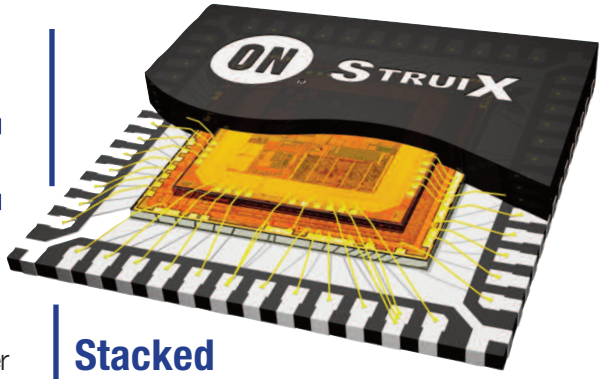

Stacked sensing

Using advanced die stacking technology, ON Semiconductor has developed a semi-customisable sensor in a System-in-Package format, targeting medical applications such as glucose monitors and electrocardiogram systems.

Named Struix, which is Latin for 'stacked', it integrates a custom-designed analog frontend alongside a ULPMC10 32-bit MCU, which is based on a Cortex-M3. A 12-bit ADC, real-time clock, PLL and temperature sensor are also integrated .

Get Smart!

Addressing the potentially damaging fragmentation in the Smart Energy market, Murata has joined the EEBus initiative, which has the goal of joining up networking standards and energy management in the smart grid.

Murata has already developed a gateway that uses wireless connectivity, which will support the EEBus standards. Murata's Rui Ramalho stated that the smart home will rely not only on 'agreed standards, but also collaboration by appliance equipment manufacturers'. Existing members include Schnieder Electric, Bosch and SMA.

Wi-Fi for the People

The Wireless Broadband Alliance (WBA) is intent on community Wi-Fi deployment, which would enable residential gateways to be opened up to casual users. Following nine months of collaboration with over 20 major providers, it has published a White Paper (www.wballiance.com/resource-center/white-papers), which will now be presented to a number of industry forums for review.

The Paper addresses the key challenges and technology gaps, as well as clarifying the business benefits of a Community Wi-Fi service, with particularly emphasis on security and network management.

Adding Connectivity to Your Design

Microchip offers support for a variety of wired and wireless communication protocols, including peripheral devices and solutions that are integrated with a PIC® Microcontroller (MCU) or dsPIC® Digital Signal Controller (DSC).

Microchip's Solutions include:

USB
8-, 16- and 32-bit USB MCUs for basic, low-cost applications to complex and highly integrated systems along with free license software libraries including support for USB device, host, and On-The-Go.

Ethernet
PIC MCUs with integrated 10/100 Ethernet MAC, standalone Ethernet controllers and EUI - 48™/EUI - 64™ enabled MAC address chips.

CAN
8-, 16- and 32-bit MCUs and 16-bit DSCs with integrated CAN, stand alone CAN controllers, CAN I/O expanders and CAN transceivers.

LIN
LIN Bus Master Nodes as well as LIN Bus Slave Nodes for 8-, 16- and 32-bit PIC MCUs and 16-bit dsPIC DSCs. The physical layer connection is supported by CAN and LIN transceivers.

Wi-Fi®
Innovative wireless chips and modules allowing a wide range of devices to connect to the Internet. Embedded IEEE Std 802.11 Wi-Fi transceiver modules and free TCP/IP stacks.

ZigBee®
Certified ZigBee Compliant Platform (ZCP) for the ZigBee PRO, ZigBee RF4CE and ZigBee 2006 protocol stacks. Microchip's solutions consist of transceiver products, PIC18, PIC24 and PIC32 MCU and dsPIC DSC families, and certified firmware protocol stacks.

MiWi™
MiWi and MiWi P2P are free proprietary protocol stacks developed by Microchip for short-range wireless networking applications based on the IEEE 802.15.4™ WPAN specification.

BEFORE YOUR NEXT WIRED OR WIRELESS DESIGN:

1. Download free software libraries
2. Find a low-cost development tool
3. Order samples

www.microchip.com/usb
www.microchip.com/ethernet
www.microchip.com/can
www.microchip.com/lin
www.microchip.com/wireless

Wi-Fi G Demo Board
(DV102412)

MICROCHIP

Microcontrollers • Digital Signal Controllers • Analog • Memory • Wireless

The Microchip name and logo, the Microchip logo, dsPIC, MPLAB and PIC are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
©2012 Microchip Technology Inc. All rights reserved. ME1023BEng/05.13

Out of the Blue

It may be late to market, but Mouser has officially launched MultiSIM Blue, a free 'cut down' version of National Instrument's simulation tool, MultiSIM. It means Mouser can now offer a free schematic capture and PCB design tool, like other catalogue distributors such as RS and Digikey.
However, uniquely, MultiSIM Blue also integrates Spice-based simulation for analogue designs, something other free tools currently don't offer; a feature enabled by NI's technology. MultiSIM is available in several versions, all paid for, which integrate varying levels of capability and features such as virtual instruments that can be placed on a schematic and used to provide and monitor signals. Unlike existing versions of MultiSIM this version, which has taken around a year of collaborative design between Mouser and NI, is free and so fulfils strategic objectives for both companies.
Perhaps more significant than the simulation capability is a feature that links the Bill of Materials in the schematic design directly to Mouser's e-commerce site, allowing components to be purchased with a few click of a mouse. With an installed user-base of around 10,000 engineers using MultiSIM globally, it offers Mouser a potentially valuable route to new purchasers.

Quick as a flash

Boasting instant-on and dual-configuration, Altera has revealed the latest addition to its Generation 10 portfolio, the MAX 10. It effectively marks the end of the family being marketed as a CPLD and, instead, labelled as an FPGA; the direction in which the architecture has evolved.

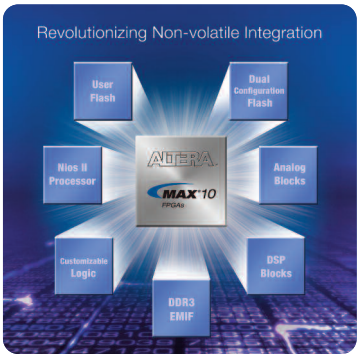

According to Altera, the device will consume power that positions it between an MCU and a traditional FPGA; 10s of mW. Using non-volatile Flash memory means the device features an 'instant-on' power-up configuration of less than 10ms and can be reconfigured from its on-board dual-configuration memory in about the same time. It's expected that the dual-configuration could be used to provide a 'safe' configuration alongside a field-upgradeable option.

V2X steps closer

Vehicle-toVehicle and Vehicle-to-Infrastructure chipsets are going in to high-volume manufacturer following a deign-win by NXP to supply Delphi Automotive. It will be the first platform to enter the market and is expected to be in vehicles within two years.

The wireless technology operates under the IEEE 802.11p standard, designed specifically for the automotive market, in order to provide real-time safety information to drivers and vehicles, delivered by a wider infrastructure comprising other vehicles, traffic lights and signage.

MEMS coverage

Shipments of MEMS sensors and actuators made by STMicro has topped 8 billion, 5 billion of which were sensor shipments, the company has disclosed. The range of applications goes beyond the smart phone to include medical, automotive and safety products, as well as others.
The hype around the IoT and now wearable technology is set to see the number of applications expand further, according to IHS Senior Principal Analyst for MEMS Sensors, Jérémie Bouchaud.

SSDs hit new heights

Mass production of 3.2Tbyte solid state drives (SSDs) has begun, doubling Samsung's previous capacity of 1.6Tbyte on a PCIe card. It uses Samsung's proprietary 3D vertical NAND memory which can sustain a sequential read speed of 3000Mbyte/s and writes at up to 2,200Mbyte/s.
Samsung plans to introduce V-NAND based SSDs with even higher performance, density and reliability in the future, targeting high-end enterprise servers.

Performance is its middle name.

KEYSIGHT M9393A PXIe PERFORMANCE VECTOR SIGNAL ANALYZER

Introducing the world's fastest, 27 GHz high performance PXIe vector signal analyzer (VSA), the realization of Keysight's microwave measurement expertise in PXI. The M9393A integrates core signal analysis capabilities and proven measurement software with modular hardware speed and accuracy. So you can tailor your system to fit specific needs today and tomorrow. Deploy the M9393A and acquire the performance edge in PXI.

M9393A PXIe performance VSA	
Frequency Range	9 kHz to 27 GHz
Analysis Bandwidth	Up to 160 MHz
Frequency Tuning	150 µs
Amplitude Accuracy	± 0.15 dB

Download new app note on innovative techniques for noise, image and spur suppression.
www.keysight.com/find/M9393AMW

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

© Keysight Technologies, Inc. 2014

Agilent's Electronic Measurement Group has become **Keysight Technologies.**

Era of the ASSoC

Swinging in favour of the custom SoC.
By Diya Soubra, Product Manager at ARM

Today, System-on-chips (SoCs) enable and underpin many applications, because a key factor in these applications is cost; by replacing several individual components with one device, the systems manufacturer will save on the bill of materials (BoM). Having fewer components also increases reliability and the mean-time between failure, alongside reducing test time and the number of defects due to PCB assembly issues.

Size and weight are also reduced by packaging components together to fit in a smaller space. Bringing functions together onto one die also leads to power and performance improvements. For vehicle manufacturers, SoCs improve control algorithms but also reduce cabling costs, as they use light, low-cost twisted-pair networks rather than the heavier cables needed to relay analogue signals.

A custom SoC provides added design security. Overbuilding becomes difficult as the contract manufacturer is able to use only the parts shipped to it by the customer, and counterfeiting is made significantly more expensive as it involves reverse-engineering the silicon. It is also possible to disguise the operation of key components, or to add product keys to each device to uniquely identify and lock it.

Advancing technologies are also making mature process nodes ideal for specialised applications. Gartner reports foundry wafer prices for a given process node decreased by an average of 10 per cent per year over the past decade, making SoC products more attractive than combining individual components. Investment in these mature nodes has also allowed startups utilising the fabless business model to be more cost-effective, as they are able to leverage the power of customisation to deliver products that offer higher performance, lower power and lower cost to customers. Cost savings mean businesses are able to go to design houses such as S3 Group, create custom SoCs and still make a healthy profit.

ARM cores are at the heart of embedded processing. The ARM Cortex-M processor family has proven successful as the basis for a wide range of microcontrollers and SoCs because it was designed for power and area efficiency. The Cortex-M0+ offers significantly more performance with higher code density compared to the 8bit architectures. The Cortex-M4 adds DSP instructions and support for floating-point arithmetic which greatly enhances the performance of sensor-driven designs. The EDA tools used to synthesise and lay out the circuits inside these SoCs are ARM-optimised and have improved dramatically in recent years, reducing the design complexity of such mixed signal devices. These techniques allow system-level optimisations for power, accuracy and performance that would be impossible using off-the-shelf parts.

The result is an environment where companies benefit from the experience of different teams to create highly differentiated, well-protected product lines that take full advantage of these highly accessible process nodes and sophisticated design tools. The custom SoC is now the smarter choice.

Ultra-Low $R_{DS(on)}$ Automotive COOLiRFET™

Package	Voltage (V)	$R_{DS(on)}$ Max@ 10Vgs (mΩ)	Q_G Typ (nc)	I_D Max (A)	RthjC Max	Part Number
D²PAK-7P	40	0.75	305	240	0.40°C/W	AUIRFS8409-7P
		1.0	210	240	0.51°C/W	AUIRFS8408-7P
		1.3	150	240	0.65°C/W	AUIRFS8407-7P
D²PAK	40	1.2	300	195	0.40°C/W	AUIRFS8409
		1.6	216	195	0.51°C/W	AUIRFS8408
		1.8	150	195	0.65°C/W	AUIRFS8407
		2.3	107	120	0.92°C/W	AUIRFS8405
		3.3	62	120	1.52°C/W	AUIRFS8403
TO-262	40	1.2	300	195	0.40°C/W	AUIRFSL8409
		1.6	216	195	0.51°C/W	AUIRFSL8408
		1.8	150	195	0.65°C/W	AUIRFSL8407
		2.3	107	120	0.92°C/W	AUIRFSL8405
		3.3	62	120	1.52°C/W	AUIRFSL8403
TO-220	40	1.3	300	195	0.40 °C/W	AUIRFB8409
		2.0	150	195	0.65 °C/W	AUIRFB8407
		2.5	107	120	0.92 °C/W	AUIRFB8405
DPAK	40	1.98	103	100	0.92 °C/W	AUIRFR8405
		3.1	66	100	1.52 °C/W	AUIRFR8403
		4.25	42	100	1.90 °C/W	AUIRFR8401
IPAK	40	1.98	103	100	0.92 °C/W	AUIRFU8405
		3.1	66	100	1.52 °C/W	AUIRFU8403
		4.25	42	100	1.90 °C/W	AUIRFU8401
5x6 PQFN	40	3.3	65	95	1.60 °C/W	AUIRFN8403
		4.6	44	84	2.40 °C/W	AUIRFN8401
5x6 PQFN Dual	40	5.9	40	50	3.00 °C/W	AUIRFN8459
		10.0	22	43	4.40 °C/W	AUIRFN8458

The new International Rectifier AEC-Q101 qualified COOLiRFET™ technology sets a new benchmark with its ultra-low $R_{DS(on)}$. The advanced silicon trench technology has been developed specifically for the needs of automotive heavy load applications offering system level benefits as a result of superior $R_{DS(on)}$, robust avalanche performance and a wide range of packaging options.

The COOLiRFET™ Advantage:

- Benchmark $R_{DS(on)}$
- AEC Q101 qualified
- High current capability
- Robust avalanche capability

Key Applications:

- Electric power steering
- Battery switch
- Pumps
- Actuators
- Fans
- Heavy load applications

for more information call +49 (0) 6102 884 311 or visit us at www.irf.com

International IOR Rectifier
THE POWER MANAGEMENT LEADER

Visit us!
Hall A5
Booth 320

electronica 2014
inside tomorrow
Messe München, November 11–14, 2014

Unseen Energy

The promise of truly untethered charging may be nearer than we think, allowing devices to recharge over greater distances, as Philip Ling discovers

With smart phones and other mobile devices about to be joined by a diverse array of remote sensors and wearable technology, keeping technology charged is becoming evermore troublesome. Energy harvesting provides part of the solution but for technology that needs to be 'always on' this represents a challenge. Battery technology continues to evolve but, arguably, what the industry really needs is a revolutionary — rather than evolutionary — step forward.

While wireless power transfer is fundamental to today's way of life, the emergence of wireless charging embodied by closely coupled inductance has emerged as a potential solution. But technologies such as Qi still require the device to have a modified battery (or power pack) and a corresponding mat, upon which the device is placed and charged through inductive coupling. While this is undoubtedly building momentum, the concept of truly wire-free charging is even more enticing; it promises that, instead of placing our devices in a charging area, they can be charged just by being in the same room as the transmitter. Although limited in power transfer capability, today's devices are typified by their ultra-low power requirements, so is the technology on the cusp of wide scale adoption?

Image courtesy of DigitalArt at FreeDigitalPhotos.net

Essentially the concept is that devices can be charged or powered while they are still in use; without the need to be plugged in, placed on or otherwise moved. Instead, power is transferred over distances measured in whole metres rather than a fraction. There is now a number of companies pioneering this technology, taking a range of approaches; UBeam uses ultrasound to transmit energy across distances, while others use more conventional wireless technology, employing vast arrays of antennas each transmitting small amounts of energy in a focused way to a single receiver. This allows relatively high levels of power to be transmitted — in some cases to several devices at once — without user intervention.

No more cables?

One such solution, expected to be demonstrated at CES in January 2015, comes from Energous; while it is still not giving away too many details, its CTO, Director and Founder, Michael Leabman, explained the fundamentals. He refers to the technology as 'wire-free' as opposed to wireless; differentiating it from solutions such as Qi which rely on small (or effectively no) distances between the charger and the device.

This means it will work up to distances of 15

feet (around 5 metres), effectively covering an entire room, although that will depend on the device's power requirements and the number of devices being charged from a single charger. For example, Leabman explained that the first solution will charge four devices at 4W at 5ft, or four devices at 2W over 5 to 10ft, or four devices at 1W at up to 15ft.

A single charging base station will be able to charge up to 24 devices at once, based on the total power requirements. This could, in theory, allow many IoT sensor nodes in close proximity to be powered by a single mains-connected transmitter, as power can be delivered continuously (going beyond 24 will require some form of time-division distribution). It achieves this by creating multiple 'pockets' of energy at 1/4W 'intervals', but it isn't simply through using one antenna per 1/4W of power: "We actually use all of our array to focus the energy on every device, that enables us to have a very focused pocket; it's probably different to what people think and it's hard to do, but it gives us an advantage and the best efficiency." Power is delivered using the same unlicensed bandwidths used by Wi-Fi and the company is targeting devices that require 10W or less.

Another company with a similar approach is Ossia, which describes itself as an 'early stage, technology-mature' company. Its solution also uses unlicensed bandwidths but goes further to explain that power will actually be received through the same antennas used in the devices for Wi-Fi or Bluetooth, over distances of up to 30ft. It's also enabled by 'smart antenna' technology and the company claims that fitting it in to devices requires just a small IC (pictured).

Both technologies offer something that traditional chargers find difficult or cumbersome; charging many devices from a single point. The key to wire-free charging will therefore be in keeping the power requirements of the device low enough to benefit from the technology. While the semiconductor industry is pushing hard in this direction, the user demands are constantly rising and that spells more power. Software control is mentioned frequently by the developers of both solutions covered here, which indicates that there will be an opportunity for 'quality of service' to be added by providers. Perhaps wire-free charging will become a part of a mobile phone tariff, but there will inevitably be some form of billing required for public spaces. Security may also be an issue.

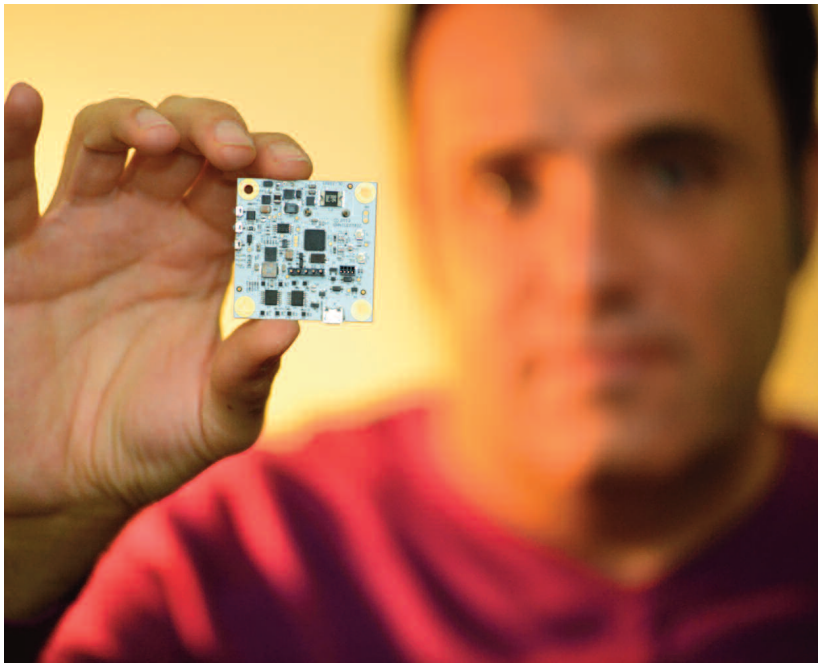

With the ink on patents for wire-free charging still drying, it's no surprise that details are sketchy. However, that doesn't stop the message getting out; wire-free charging is just around the corner (although it doesn't necessarily go around corners) and soon we will all enjoy the freedom of our devices drawing power while still in our hands, pockets, bags or backpacks.

Image courtesy of CoolDesign at FreeDigitalPhotos.net

Harnessing the Tides

An ambitious project to generate tidal energy in Scotland requires some very specialised electronics, reports Sally Ward-Foxton

A huge project to install tidal turbines to generate electricity off the coast of Scotland is set to generate 86MW by 2020. There is potential for turbines to generate 398 MW, enough to power 175,000 Scottish homes from the same tidal energy farm once the project is fully rolled out. Phase 1A of the MeyGen project, in the Pentland Firth where tidal flow can reach 5m/s, will see four 1.5MW demonstrator turbines installed under the sea to harness the power of the tides.

Tidal turbines can produce similar efficiency to wind turbines, with the MeyGen turbines expected to deliver up to 40% 'water-to-wire' overall efficiency. However, unlike wind, tides have the advantage of being predictable. This means the turbines can produce the full 1.5MW of power for a larger proportion of the time – MeyGen's load factor is expected to be around 40%, compared to a wind farm's typical load factor of 20-25%. This also means that the amount of energy, and therefore, revenue for the tide farm can be estimated very accurately over the lifetime of the turbines (25 years). The tidal turbine technology is based on today's wind farms, with some key differences to account for the sub-sea environment, explains David Taaffe, Head of Electrical Systems for the project.

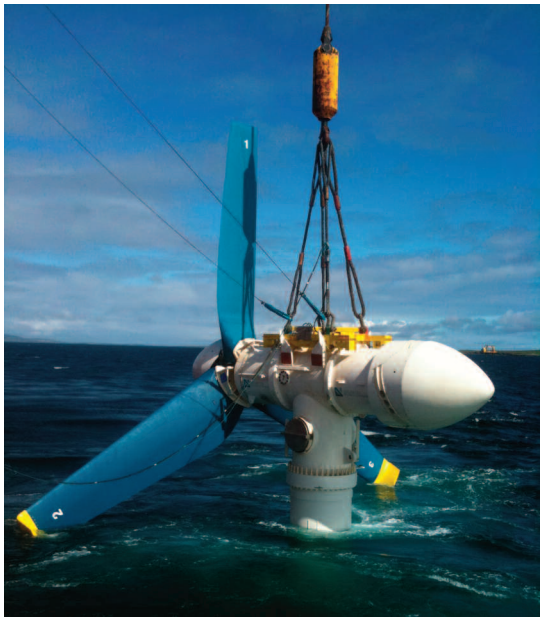

Figure 1. An AR1000 tidal turbine is deployed at the European Marine Energy Centre in Orkney. The turbine section is carefully lowered into the water from the deck of the installing vessel using its crane.

"The main components in the turbine system are the same, an electrical generator, gearbox, electrical brakes, frequency converters, transformers and cables, with sub-sea connectors. But the location has been selected so that the more vulnerable components are located on-shore," Taaffe explains. "We've split the turbine at the generator. With a wind turbine the frequency converter would be inside the nacelle, on our machines, it's on-shore, around 2.5km away."

Placing the frequency converters in an on-shore power conversion centre means they are more easily accessible for operation and maintenance purposes. Each turbine then requires a 2.5km cable to it, whose installation impacts the economics of the project. However, compared to the cost of performing sub-sea maintenance (the maintenance vessels can cost anywhere between £30,000 and £70,000 a day, according to Taaffe), it's seen as an acceptable compromise. Having separate cables to each turbine also means that each turbine is independent – if a fault was to switch off one turbine, the rest of the farm would be unaffected.

"Once we put the turbines sub-sea, we're aiming for a 5-year mean time between repairs," Taaffe says. "It's common sense that the fewer components you have in the turbine, the longer the time will be between repairs."

THE TOTAL SOLUTION

Comprehensive programs for the critical semiconductors needed... with no compromises.

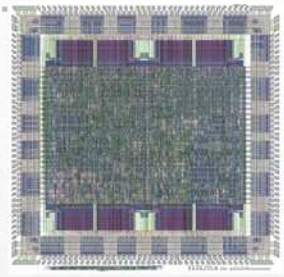

Rochester Electronics® www.rocelec.com • 978.462.9332

Re-introducing Intel's MQ80C196KB CHMOS Microcontroller

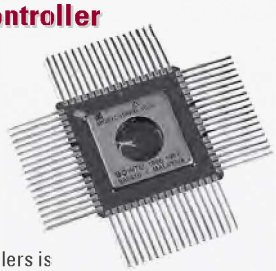

The Intel MCS-96 Family of microcontrollers is most commonly used in embedded systems for motor control, modem, printer and pattern recognition applications. The 68-pin microcontroller from Rochester is pin-for-pin compatible, uses a true superset of the M8096 instructions and is offered in a variety of packages, speeds and temperature ranges.

Datasheet

Check Inventory Here

Re-introducing Altera's EP1810GC-35 EPLD

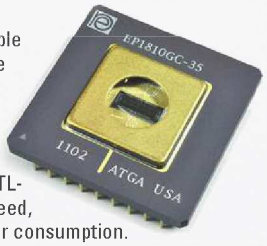

The complex EP1810 erasable programmable logic devices (EPLDs). The EP1810 devices offer LSI density, TTL-equivalent speed, and low power consumption. EP1810 devices have 48 macrocells, 16 dedicated input pins, and 48 I/O pins. EP1810 devices are divided into four quadrants, with each containing 12 macrocells. Macrocells can be individually programmed as D, T, JK or SR flipflops, or combinatorial operation. Available in a 68-pin ceramic PGA package or a 68-pin J-lead package, the EP1810.

Datasheet

Check Inventory Here

Re-introducing AMD's AM27C512 CMOS EPROM

The AM27C512 product family is a 512-Kb, UV erasable PROM, organized as 64K words by 8 bits and operates from a single 5V supply. Created for high speed and low power, the 27C512 is available in windowed ceramic DIP packages as well as one-time programmable PDIP and PLCC package.

Datasheet

Check Inventory Here

Figure 2. A series of guide wires are used to control the orientation of the nacelle, enabling the accurate alignment of the nacelle with the base substructure.

The turbine's frequency converter has two functions. Firstly, it's a rectifier for the power coming from the generator in the turbine, which delivers various power levels at variable frequencies (between 200kW and 1.5MW, and between 10 and 50Hz, depending on prevailing tidal conditions). Rectifying it to DC, it can then be easily reconverted to 50Hz, 33kV so it can be linked to the grid from there. The medium-voltage frequency converter runs at 3.3kV in order to overcome losses in the 2.5km cable. Secondly, the frequency converter has an inverter which is used to drive the system which controls the torque coming from the generator.

"Because of the nature of tidal flows and water turbulence, each turbine has to be kept tightly controlled to stop it from going over speed or vibrating too much," Taaffe says. Torque control is required to limit the power output of each turbine to 1.5MW, because that's the level to which the components are rated. At the same time, the generators have an optimum RPM in order to extract maximum power from the tide. Torque control is achieved through adjusting the pitch of the turbine blades. The system is again based on proven wind farm technology, but with modifications made for the under-sea environment.

"The pitch controller has a PLC that takes information from various parts of the system to decide what pitch angle to move to, and then it has a closed loop controller which takes feedback following each decision on what the power levels are, and what the next angle needs to be," Taaffe says. "Eventually, it will regulate the power… if the current power levels are higher than the desired power levels, it will pitch out of the flow to decrease the power. If the tide then drops, it will overshoot, so it will pitch back in. At high flows, you'll have continuous operation of the blade pitch system."

Clearly, the algorithm for wind turbines is not going to work under the sea, because it needs to be tuned for the inertia of the system and for water dynamics as opposed to air dynamics. Factors like the time it takes to adjust the pitch and the acceleration of the blades all needed to be tuned in the real environment, and that has been achieved at the European Marine Energy Centre in Orkney, where a number of turbines have been tested and proven with the pitch controller. Figures 1 and 2 show a demonstrator being installed at EMEC, with Figures 3 and 4 showing an illustration of the same turbine during installation and in place on the sea bed.

Figure 3. The turbine section is lowered directly onto the already installed base section using underwater cameras for guidance. The whole thing is weighted down to the sea bed with large ballast. Each of the three pieces of ballast weighs 200 tonnes, as that is the limit of the transporting equipment.

Figure 4. The turbine rests on the base using a simple gravity-based mechanism, which requires no bolts or clamps. It's ready for operation within 60 minutes of leaving the deck of the vessel.

The demonstration phase of the MeyGen project will actually use two completely different tidal turbine designs, which will be evaluated. Three turbines will be supplied by Andritz Hydro Hammerfest, the HS1000 model, and one turbine will be supplied by Atlantis, the AR1000 (Figure 5). Both turbine types have been designed to fit the design envelope shown in Figure 6, but they have very different electrical systems. Andritz Hydro Hammerfest's design for the HS1000 uses an induction generator, while the AR1000 from Atlantis uses a permanent magnet generator. Each has pros and cons, explains Taaffe.

"The permanent magnet generator is more efficient, the induction machine is a bit more lossy, but the permanent magnet machine is probably slightly less robust than the induction machine," he says.

The superior efficiency of the permanent magnet generator design is down to the fact that for a given power, the permanent magnet generator is smaller than the induction machine. This means that the permanent magnet machine can be made larger without compromising the constraints on the diameter of the nacelle, so it can be run at slower speeds. This means the permanent magnet generator needs only a two-stage gearbox, compared to the induction generator's three-stage gearbox, improving the overall efficiency of the design.

Andritz's induction generator design uses a standard

NEW LabVIEW 2014 NOW AVAILABLE!

Infinite Designs, One Platform

with the only complete system design environment

LabVIEW is the only comprehensive development environment with the unprecedented hardware integration and wide-ranging compatibility you need to meet any measurement and control application challenge. LabVIEW is at the heart of the graphical system design approach, which uses an open platform of productive software and reconfigurable hardware to accelerate the development of your system.

>> Accelerate your system design productivity at ni.com/labview-platform

Follow us on *Search* **National Instruments** *or* **LabVIEW**

NATIONAL INSTRUMENTS

Austria 43 662 457990 0 ▪ **Belgium** 32 (0) 2 757 0020
Czech Republic, Slovakia 420 224 235 774 ▪ **Denmark** 45 45 76 26 00
Finland 358 (0) 9 725 72511 ▪ **France** 33 (0) 8 20 20 0414 ▪ **Germany** 49 89 7413130
Hungary 36 23 448 900 ▪ **Ireland** 353 (0) 1867 4374 ▪ **Israel** 972 3 6393737
Italy 39 02 41309277 ▪ **Netherlands** 31 (0) 348 433 466 ▪ **Norway** 47 (0) 66 90 76 60
Poland 48 22 328 90 10 ▪ **Portugal** 351 210 311 210 ▪ **Russia** 7 495 783 6851
Slovenia, Croatia, Bosnia and Herzegovina, Serbia, Montenegro, Macedonia 386 3 425 42 00
Spain 34 (91) 640 0085 ▪ **Sweden** 46 (0) 8 587 895 00 ▪ **Switzerland** 41 56 2005151
UK 44 (0) 1635 517300

©2014 National Instruments. All rights reserved. LabVIEW, National Instruments, NI and ni.com are trademarks of National Instruments. Other product and company names listed are trademarks or trade names of their respective companies. 07926

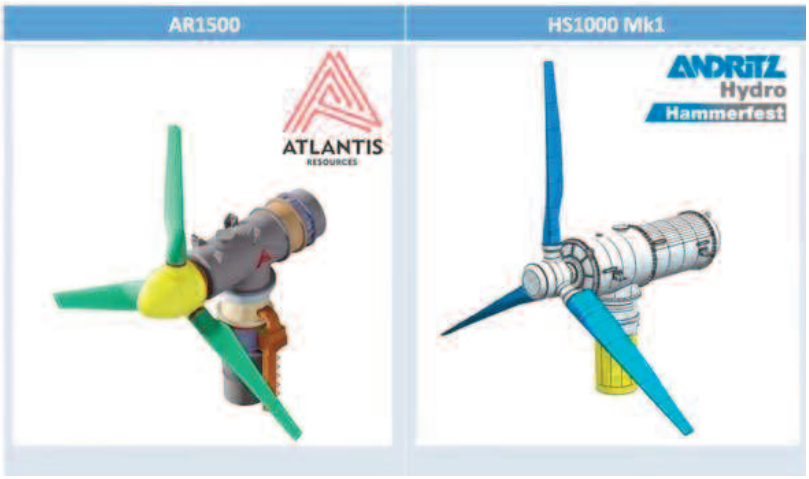

Figure 5. The two types of turbine chosen for Phase 1A of the project are the Atlantis AR1000 and the Andritz Hydro Hammerfest HS1000.

off-the-shelf 1500 rpm induction machine, with a standard three-stage gearbox, like a wind turbine. It's made waterproof by putting it inside a thick steel housing. "Although the steel casing is an additional cost, it's actually a very conservative approach," Taaffe points out. "With waterproofing standard components, you've got a lot of confidence that the design will work, because it's essentially a wind turbine in a waterproof pod."

The waterproof nacelle of this design has air inside, and the components are cooled with oil that travels down pipes to an external heat exchanger, which dissipates the heat generated into the sea water. Testing has proven that actively cooling the major components in the system works well, but because it uses electronics in such an inaccessible location, it does require redundancy in terms of additional power supplies and controllers.

Atlantis's permanent magnet generator design's components are designed to be waterproof in their own right, with a special seal between the gearbox and the generator; with the thick steel waterproof casing not required, the nacelle can have a smaller diameter. "This approach is a bit more novel," says Taaffe. "The components are a bit more bespoke, and the mechanical arrangement is more bespoke, so it requires more testing before it can be installed sub-sea."

Atlantis's design is going to be passively cooled; since the nacelle is smaller and the case thinner, the components are closer to the seawater passing over the structure. The case is filled with oil, and when the shaft rotates, it circulates the oil around the structure, allowing heat to be transferred from the magnets and coil to the oil, then to the outer casing and out into the seawater. Taaffe points out that a key benefit for tidal systems is that the highest power levels, and therefore the highest levels of heat to be dissipated, occur when the water is flowing fastest, so it lends itself to passive cooling.

The blade pitch systems are also different. In Andritz's design, the pitch of each blade is controlled individually, with each blade having its own motor and gears. Electronics is used to control the pitch angle, with each blade having a frequency converter to drive the motor (in fact, there are two frequency controllers per blade, for redundancy). The Atlantis design has a collective pitching system whose mechanical gears pitch all three blades at the same time. This system is hydraulic, with an electric motor to repressurise the hydraulics.

Installation of the four turbines for Phase 1A of the project in the Pentland Firth will commence in 2015. If everything goes to plan, the completion of Phase 1 (parts A, B, C and D) will result in 57 turbine installations giving the wind farm a capacity of 86MW by the end of 2020.

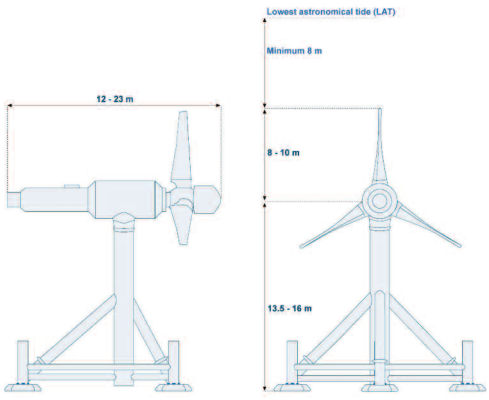

Figure 6. The design envelope for the turbines, showing the major dimensions.

Convection-cooled Medical Power with Ultra High Efficiency

XP Power provides a wide range of medically approved power supplies with convection-cooled ratings for all healthcare power solutions.

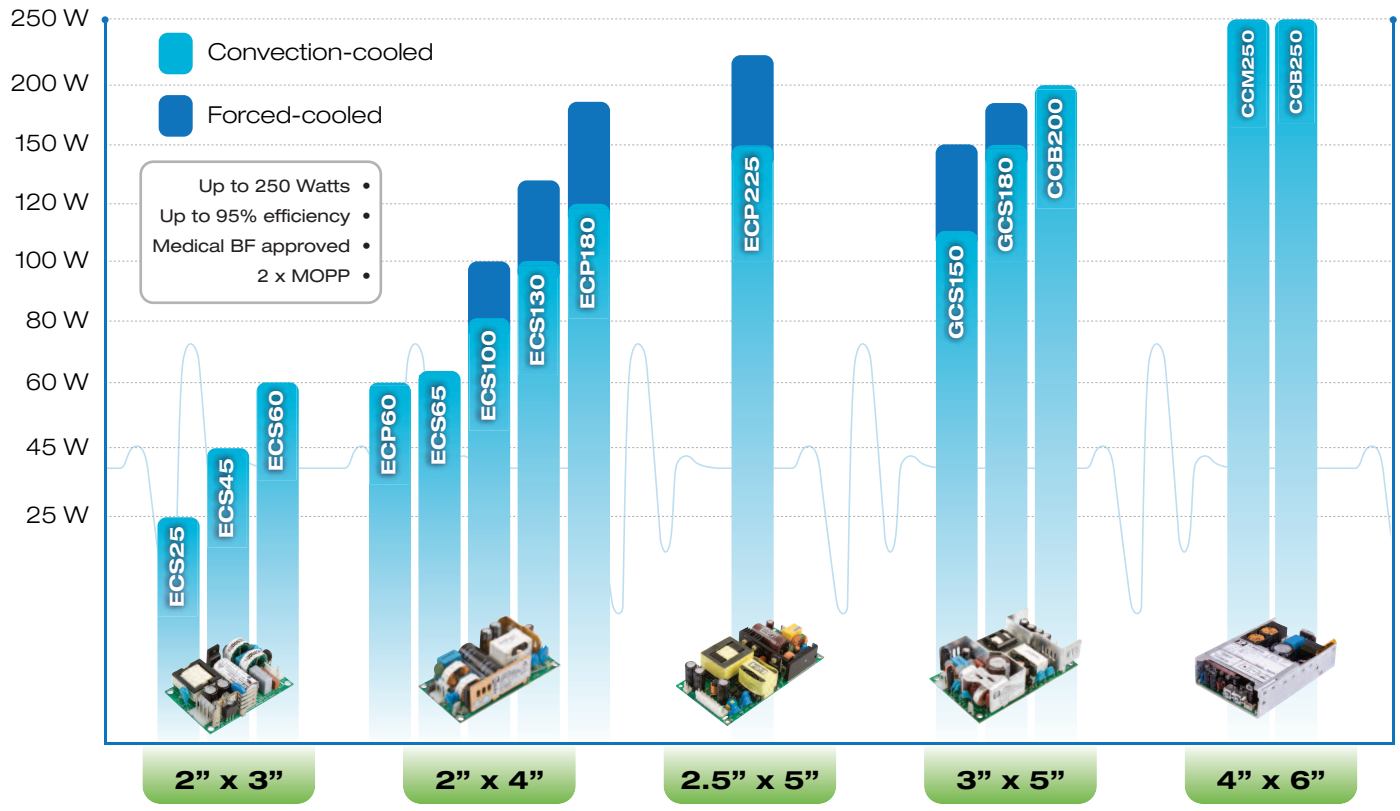

Visit our website for more information or to request a copy of our new Power Supply Guide and see our complete line of power products.

GREEN XP POWER

Selector App Available

XP Power
www.xppower.com

Charging Towards Fewer Cables

Wireless charging provides a convenient way for users to charge all their mobile devices simultaneously, as a result deployment is growing. By David Zelkha, Managing Director of Luso Electronics

There is nothing new about the concept of inductive charging, also known as wireless charging. Electric toothbrush users have been familiar with it for many years; but the proliferation of smartphones and tablets, and the associated difficulties with the various charging connectors involved, more companies are taking an interest in the idea.

The physics of wireless charging is straightforward; power is transferred between two coils using inductive coupling. The charging unit contains one of the coils that acts as the transmitter and the receiver coil is in the unit to be charged. An alternating current in the transmitting coil generates a magnetic field which induces a voltage in the receiving coil, and that voltage can be used to power a mobile device or charge a battery.

Palm was one of the pioneers of wireless charging for mobile phones using what were, at the time, innovative coil sets from E&E Magnetic Products (EEMP) and just a few years ago the problem of charging these devices, with their energy sapping Wifi and Bluetooth features, was widespread with almost every phone having a different connector. The advent of mini-USB, used by many of the mobile phone makers, has eased this problem somewhat but the convenience of wireless charging is attracting growing attention.

However, its use has been held back due to a couple of factors. One was the lack of a standard, but that has been solved, more of which later. And the other related problem was unwillingness by most phone makers to install the necessary coil in a device that is already crammed with electronics. Some companies introduced sleeves that could be put over the phone but that in a way lost the convenience of just being able to drop the phone on a charging mat.

A change in the market happened in 2009. Before that, the reluctance to put the coil into the phone was understandable given that it wasn't just the coil but a number of other components that went with it. But this altered when Texas Instruments introduced the first chipset for this. It was followed soon by IDT and now there are five or six companies – such as NXP, Freescale, Panasonic, Active-Semi and Toshiba –

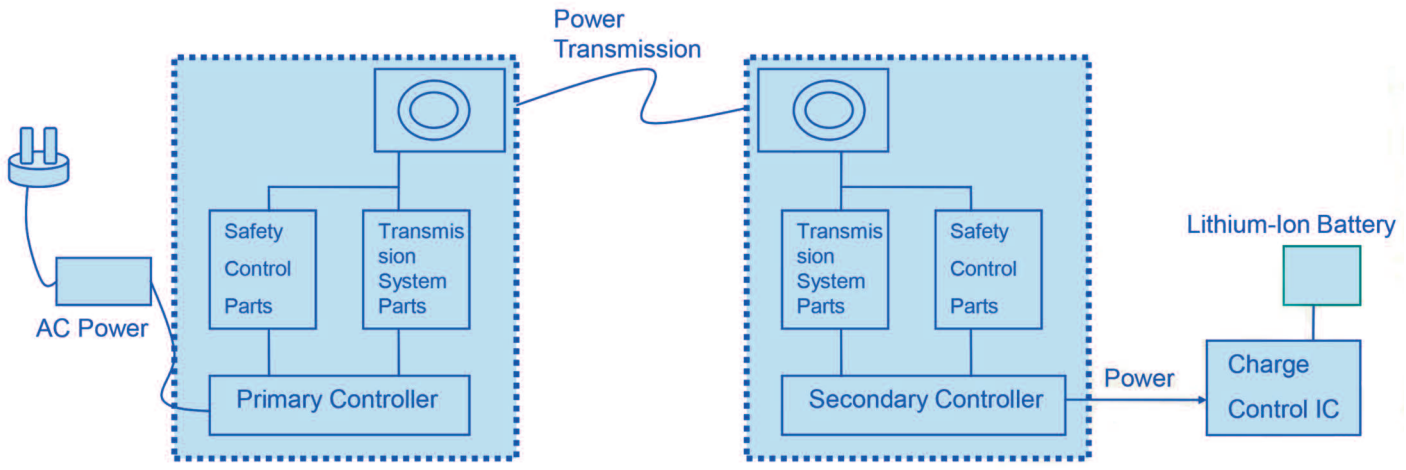

Figure 1: Basic overview of wireless charging technology

making chipsets, and this has helped drive interest from the phone makers.

Leading the charge

Some car makers are also looking at this technology to provide drivers and passengers with an easy way to charge their devices while in the car. Toyota has already brought out models with this installed. The 2013 Toyota Avalon limited edition was the first car in the world to offer wireless charging based on the Qi standard. The in-console wireless charging for Qi–enabled devices was part of a technology package that included dynamic radar cruise control, automatic high beams and a pre-collision system.

The Avalon's wireless charging pad was integrated into the lid situated in the vehicle's centre console. The system can be enabled by a switch beneath the lid, and charging is as simple as placing the phone on the lid's high-friction surface. Cadillac has announced that it will have wireless charging in the 2015 model of the ATS sport sedan and coupe. It will also add the technology to the CTS sports sedan in autumn 2014 and the Escalade SUV at the end of 2014. Mercedes Benz has adopted the Qi standard for its wireless charging plans. Most other car makers are also looking at the technology. And aftermarket charging mats that can be fitted in most vehicles are available from numerous suppliers.

Once a large number of vehicles have this installed, then it is another driver to the phone manufacturers as it shows there is an available market and that this is a desirable feature. Also, once phone users experience the convenience of wireless charging in their car, they are likely to drive demand for this to be available at home and in the workplace as well.

The main limitation on wireless charging is the short distance needed between the transmit and receive coils, which must be close enough to ensure a good coupling. The technology works by creating an alternating magnetic field and converting that flux into a current in the receiver coil. However, only part of that flux reaches the receiver coil and the greater the distance the smaller the part.

The higher the coupling factor the better the transfer efficiency. There are also lower losses and less heating with a high coupling factor. When a large distance between the two coils necessary it is referred to as a loosely coupled system, which has the disadvantages of less efficiency and higher electromagnetic emissions, making in unsuitable for some applications.

A tightly coupled system is where the transmit and receive coils are the same size and the distance between the coils is much less than the diameter of the coils. Given that people are used to using charging cables that seem to get shorter with each new phone, this is not seen as a major issue. Charging time is roughly the same as with a traditional plug-in charger. Also, as most of the charging is done by placing the phone directly on a mat that contains the charger, there is no significant distance involved.

Standards

There are three main standards covering inductive charging, with the basis of all three being the same. The Alliance for Wireless Power (A4WP) standard has a higher switching frequency, which allows a greater charging distance. The Power Matters Alliance (PMA) Powermat standard also covers the alternative resonant charging method. There is talk that these two standards could soon be merged into one.

But WPC (Wireless Power Consortium) is the oldest and the most adopted standard in the market, with some 1200 different products in the market. The consortium has more than 200 members and its standard is known as Qi (pronounced 'chee'), which specifies the whole charging circuit on the transmit and receive side and how it should be implemented in the charging devices.

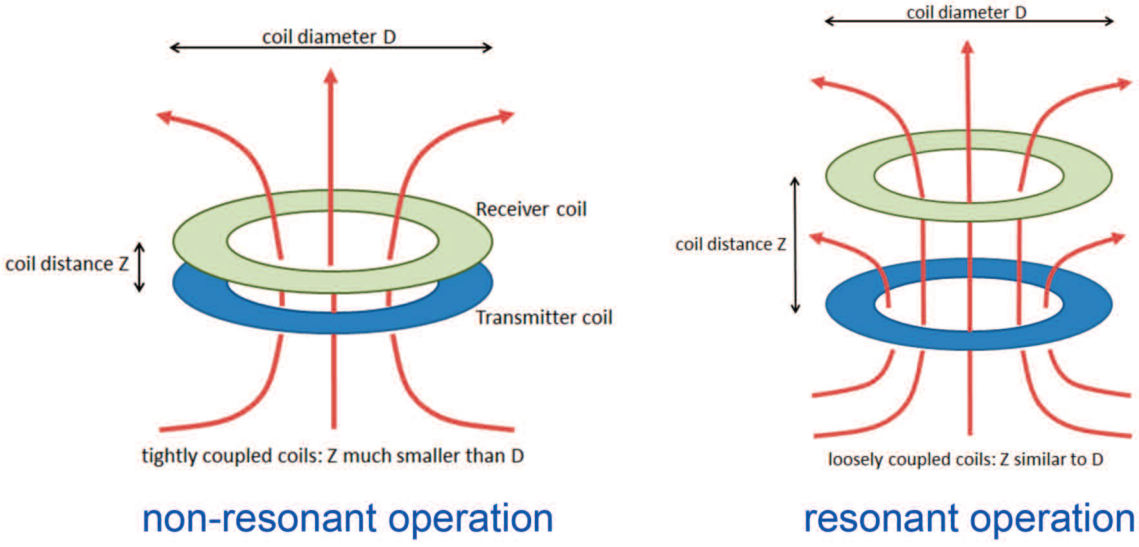

Figure 2: The difference between resonant and non-resonant operation

Most Qi transmitters use tight coupling between the coils and operate the transmitter at a frequency that is slightly different from the resonance frequency. Even though resonance can improve power transfer efficiency, especially for loosely coupled coils, two tightly coupled coils cannot both be in resonance at the same time. The Qi standard therefore uses off-resonance operation because this gives the highest amount of power at the best efficiency. However, there are Qi-Approved transmitters that will operate at longer distances, loosely coupled and at resonance.

This shows that Qi is an evolving standard. As new applications and requirements come along, the WPC is adding to the standard to keep it up to date. For example, the original standard had just one transmit coil and one receive coil. This was then extended to three transit coils to give a greater area on which a device being charged could be placed. And some now have even more coils; the standard already covers up to five. This makes it easier for users dropping the phone onto the mat, as they no longer have to position it carefully. Tightly coupled coils are sensitive to misalignment, but a multi-coil mat can be used to charge more than one device at the same time.

This is seen as one of the big advantages of wireless charging; users need just one mat onto which they can put all their phones, tablets and cameras and have them charged simultaneously. As well as for home or office use, this adds convenience for those who travel and are regularly staying at hotels – only one charger needs to be packed. There is also now a number of charging areas in public places. These have been installed in airports in Asia and the USA. Japan has more than 3300 public locations where consumers can charge their devices wirelessly. And even the French Open tennis tournament had Qi chargers in the guest areas.

There is an environmental benefit as well. The current system sees many corded chargers thrown away each time users upgrade to new mobile devices. A wireless mat that complies with the relevant standard can carry on being used with the new device if it too complies. And multi-coil transmitters allow the power to scale with increasing power levels by powering more coils underneath the receiver. The first smart phones needed 3W, whereas today's devices require over 7.5W and this is growing. Tablets, e-book readers and ultrabooks need from 10 to 30W. A loosely coupled system can achieve multi-device charging with a single transmitter coil, provided it is much larger than the receiver coils and provided the receivers can tune themselves independently to the frequency of the single transmitter coil.

In the early days, there were also problems if, say, a coin from the user's pocket was pulled

out at the same time as the phone and sat between the phone and the mat. The mat would then try to charge the coin instead. The standard was thus changed to bring in what was called foreign object detection so that the likes of coins and keys are recognised and not charged. Another problem that was sorted early on was the fear that the switching frequency could interfere with some automotive applications such as remote door opening. If a phone maker wanted a new design, then the WPC is willing to work with the company to adapt the standard to suit so the manufacturer can still use the Qi logo on the device. Compare this with A4WP, which only covers single coil, loosely coupled resonant systems.

Design-in

The main components are the chipset, the transmit and receive coils, and some passives. All the available chipsets meet the WPC standards even though packaging and pin-outs may be different. Texas Instruments and IDT are the main players in the North American market and NXP is the largest player in Europe. There are multiple suppliers for the coils, of which EEMP was the first.

Most of the research and development work is in customising these products, as there tends to be an even split in the market between standard products and customised versions. Some phone makers, for example, want really thin receive coils to keep the size of the phone down. This can cause problems as the ferrite used tends to be brittle, but EEMP has made a flexible ferrite to get round this. This makes it easier to deploy where there is a curved surface rather than moulding a standard ferrite into the correct shape. The moulded ferrite route is also more expensive.

EEMP, which is a member of the three main standard bodies – WPC, A4WP and PMA – has transmitter and receiver coil standard modules plus units that can be tailored to the size, thickness and shape needed by the application. These are available from Luso Electronics, which can provide samples and have small batch quantities available to support prototype and lower volume builds. They comply with the Qi standard and the two ranges provide 5W and the 15W extension to the Qi standard. Available in various operating temperature ranges, they are RoHS compliant, halogen free and provide low resistance and low temperature rise when operating.

The push now is to develop clever and novel techniques to expand the available market for this technology beyond smartphones and the like. Such applications could include audio systems, torches, LED candles, test equipment, handheld instruments and PoS equipment. The potential market is expanding, demonstrating the flexibility of the technology. More and more phone makers are adopting the Qi standard for wireless charging and most now have at least some models in their range that are capable of this. As car manufacturers start to introduce wireless charging in their vehicles, an increasing number of consumers will experience the convenience and start looking for this as a desirable feature when they choose their next smart phone. The result should be a growing market for this technology.

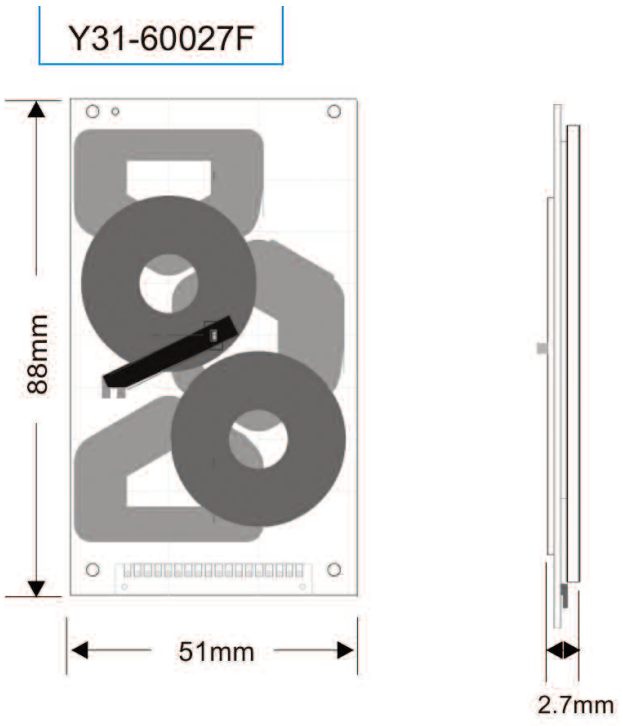

Figure 3: Multi-coil systems can use different shapes to increase the charging area, as can be seen with this configuration from EEMP

Batteries are included

Wayne Pitt, Saft's Business Development Manager for lithium batteries, explains why battery power is a key element in creating the IoT

The IoT is a hot topic right now. It promises a multitude of interconnected devices equipped with embedded sensors and intelligent decision making – storage tanks that create an alert when they need filling; household appliances that manage themselves, and bridges that monitor their own structural integrity are all examples. But the IoT concept is still a long way from reality and the underlying technologies are very much in development.

To qualify for the IoT, a device must have its own IP address and, in the industrial IoT realm, many devices will take the form of remote sensors; each made up of the sensor itself, a microprocessor and a transmitter. Some people estimate that there is scope for tens of billions of devices, many of which will be sensors embedded into the fabric of their surroundings, which feed performance data back to central databases for monitoring.

A typical sensor device will draw just a few µA in sleep mode, when it might be waiting for an external cue to take a reading. Alternatively, it might draw around 80µA in standby, running its internal clock between timed sensor readings. Data recording and processing might use 20mA and then transmitting a few tens of bytes of data might call on up to 100mA. Other devices acting as base stations or gateways will receive and relay data transmissions from many terminal devices such as environmental sensors. These applications will require higher currents and more frequent transmissions, and so will consume more energy.

With many sensors being located in hard to reach, inhospitable and remote locations, it becomes challenging to provide energy to sensors or terminal devices on the IoT. Mains power is often impractical due to the physical location or installation budget. As a result, many IoT devices will rely on batteries to provide the energy for a lifetime of operation. Replacing a battery on a sensor embedded in a high ceiling might involve extensive scaffolding, a specialist access contractor or downtime of critical equipment, with the cost of the change far outweighing the cost of the battery. So it's important to select a battery that will deliver a long and reliable life.

Battery selection

When selecting an energy source, engineers need to choose either a primary battery or a rechargeable battery operating in conjunction with some method of harvesting energy from the environment; typically a solar panel. Selection is governed by the energy requirements of the device and its application.

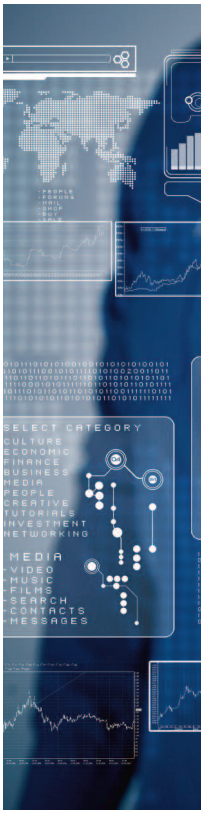

A rule of thumb is that if a device will use more energy than can be supplied from two D-sized primary batteries over a life of ten years, then a rechargeable battery is the most practical choice as it will free up the designer to use relatively high power consumption. And because not all D-sized batteries are created equal, Saft has set this threshold at between 90-120Wh (Watt-hours); the energy

stored in two LS batteries, which are based on lithium-thionyl chloride (Li-SOCl2) cell chemistry.

Finding the right battery for a new application can be extremely challenging. The power and energy requirements need to be considered against the technical performance of different battery types. Electrochemistry has a strong bearing on how a battery will perform, but other aspects such as the way a cell is constructed are also important to a battery's performance. Quality of the raw materials used in battery construction also has a major bearing on life, as do the construction techniques used on the production line. And when a battery needs to operate in a potentially extreme environment for a decade or more, the proven reliability of cells becomes a vital consideration.

The ideal primary battery for an IoT sensor has a long life, requires no maintenance, has an extremely low rate of self-discharge and delivers power reliably throughout its life, with little degradation, even towards the end of its life. In addition, because many devices will be located in harsh environments, cells should deliver current reliably in extreme temperatures.

A number of primary lithium cell chemistries are available and of these, Li-SOCl2 is currently the best fit. This is

YOU CAN'T COPY EXPERIENCE

PRECISION AND POWER RESISTORS

We invented the Manganin® resistance alloy 125 years ago. To this day, we produce the Manganin® used in our resistors by ourselves.

More than 20 years ago, we patented the use of electron-beam welding for the production of resistors, laying the foundation for the ISA-WELD® manufacturing technology (composite material of Cu-MANGANIN®-Cu). We were the first to use this method to manufacture resistors. And for a long time, we were the only ones, too.

Today, we have a wealth of expertise based on countless projects on behalf of our customers. The automotive industry's high standards were the driving force behind the continuous advancement of our BVx resistors. For years, we have also been leveraging this experience to develop successful industrial applications.

The result: resistors that provide unbeatable excellent performance, outstanding thermal characteristics and impressive value for money.

ISABELLENHÜTTE
Innovation by Tradition

Isabellenhütte Heusler GmbH & Co. KG
Eibacher Weg 3 – 5 · 35683 Dillenburg · Phone +49 (0) 2771 934-0 · Fax +49 (0) 2771 23030
sales.components@isabellenhuette.de · www.isabellenhuette.de

because it has an extremely low rate of self-discharge and a long track record over 30 years, having been widely deployed in applications such as smart metering. Li-SOCl2 cells, such as Saft's LS series, are available in the well recognised formats from ½ AA through to D, which makes them a direct mechanical replacement for conventional alkaline cells. But it should be noted that the lithium electrochemistry provides a significantly higher nominal cell voltage of 3.6V (against 1.5V). The energy cells are designed specifically for long-term applications of up to 20 years and deliver base currents in the region of a few µA with periodic pulses of up to 400mA. While the power cells can deliver pulses of an order of magnitude greater, up to 4000mA.

An important aspect is the continuity of performance throughout the life of the battery. When powering a child's toy, it doesn't matter so much if performance drops off towards the end of the battery's life but in the IoT predictable high performance can be vital. An IP enabled sensor might need to transmit data that is essential to safety or business continuity; Li-SOCl2 chemistry continues to deliver high performance throughout its life.

LS batteries are already commonly used in sensors and smart meter devices. In October 2013, Saft won two major contracts to supply the batteries to major OEMs in China for installation in gas and water meters. During the life of the meters, the batteries will provide 'fit and forget' autonomous power for a minimum 12 year service life. M2M specialist manufacturer Sensile Technologies also uses the same cell type in its SENTS smart telemetry devices for oil and gas storage tanks. The cells power the devices that measure liquid or gas tank levels, record the data and transfer it by SMS (short message service) or GPRS (general packet radio service) to a central monitoring system. Over their life of up to 10 years, the SENTS devices provide data on the levels of fuel in tanks. This data allows Sensile Technologies' customers to optimise their purchasing of fuel and other stored fluids.

Elsewhere, spiral wound LSH batteries have been selected to power telemedicine devices manufactured by SRETT. The batteries provide three years autonomous operation for the T4P device, which is designed to monitor sleep apnea patients' use of their medical devices and send performance data every 15 minutes via GPRS communication, where it can be monitored by medical professionals.

Other primary lithium battery types may also have a role to play, particularly in applications that demand high pulses of energy. An example

might be to provide relatively high power for the 'ping' of a corrosion sensor on a remote oil pipeline that uses an ultrasonic pulse to determine the thickness of pipework.

Rechargeable batteries

Lithium-ion (Li-ion) batteries make a practical choice for IoT applications that call for rechargeable batteries because of their high cycling life and reliability in extreme temperatures. There are several types of Li-ion cell chemistry, which can be blended or used individually. Of these, nickel manganese cobalt (NMC) is particularly interesting because it operates reliably across the widest temperature range.

Using NMC of Saft's own formulation, Saft can deliver a rechargeable cell that operates at temperatures between -30°C and +80° C, which means that it can provide reliable power for devices installed anywhere from an arctic blizzard to a pipeline running through a desert or integrated into equipment in an engine room. And while some Li-ion technologies suffer degradation if left on float charge (for example, consumer device batteries may degrade if left to charge continuously or for long periods), Saft NMC does not. This means that it can be paired with a solar panel and left to charge day after day without losing performance — an advantage that could represent cost savings for a sensor's operator.

A typical application is AIA who are utilising a Saft rechargeable Li-ion battery in their Solar Battery product line for Class 1 Div 2 hazardous environments and extreme temperature operation.

With a Li-ion battery operating in combination with a solar panel, AIA's Solar Battery solutions provide power, data acquisition and wireless internet communication to simplify installation, maintenance and support of remote hazardous environment sensors. The AIA system allows customers such as Taqa North to achieve end-to-end monitoring and control of their fixed or mobile assets in extreme climatic conditions.

While IoT devices are not yet commonplace, the battery technologies are already available to power them effectively and reliably, thanks to the extensive field experience built up in comparable applications such as wireless sensor networks, machine to machine applications and smart metering. Ultimately, selecting the right battery depends on having a solid understanding of the base load and pulse current that a sensor, terminal device or gateway will draw. Designers can play an important role by optimising their application, ideally keeping the size and frequency of data transmissions to a minimum. In most cases, only a handful of bytes need to be transmitted, and existing battery technology can handle this comfortably to deliver upwards of ten years of reliability in even the most demanding environments.

Mixing it up

How hybrid relays combine mechanical and solid-state technologies to deliver the best of both worlds, while making it simpler to comply with the latest legislation for energy-using devices. By Benoit Renard, SCR/Triac Application Engineer & Laurent Gonthier, SCR/Triac Application Manager, STMicroelectronics

Hybrid-relays combine both a static-relay and a mechanical-relay in parallel, marrying the low voltage-drop of a relay to the high reliability of silicon devices. Motor starters and heater controls in home appliances are already common applications, but as RoHS compliance could render mechanical relays less reliable in power switching applications they are set to become even more attractive.

It can, however, be more challenging to implement the right control for this hybrid than it might seem at first glance; voltage spikes that may occur at the transition between the mechanical switch and the silicon switch could cause electromagnetic noise emission. This article offers advice on developing a control circuit thank can reduce these voltage spikes.

When choosing an AC switch, there are well known pros and cons to selecting either mechanical or solid-state technology. The advantages of silicon technology is its faster reaction time and the absence of voltage bounces at turn-on and sparks at turn-off; a main cause of Electro Magnetic Interferences (EMI) and shorter relay life-time. The advantages of an electromechanical solution are mainly reduced conduction losses, which avoid the use of a heat-sink for applications above approximately 2A RMS, and the insulation between the driving coil and the power terminals, which renders any opto-couplers to drive silicon-controlled rectifiers (SCR) or triacs useless.

Another solution consists of using both technologies to implement a Hybrid Relay (HR) with one solid-state relay in parallel to an electromechanical relay. Figure 1 shows such a topology used in a motor-starter application. Only two hybrid relays are used here for this 3-phase motor-starter: if both relays are OFF, the motor will remain in the off-state as long as its neutral wire is not connected. An HR could also

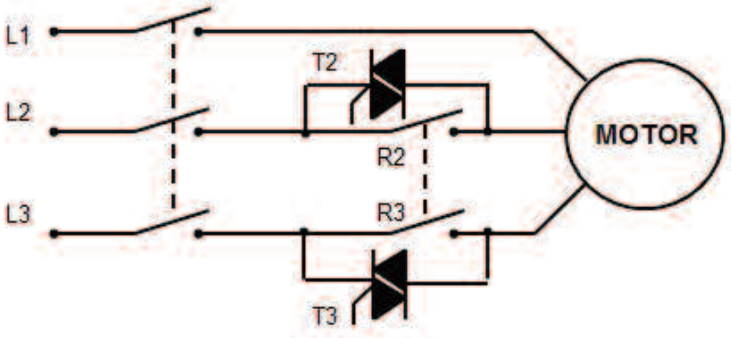

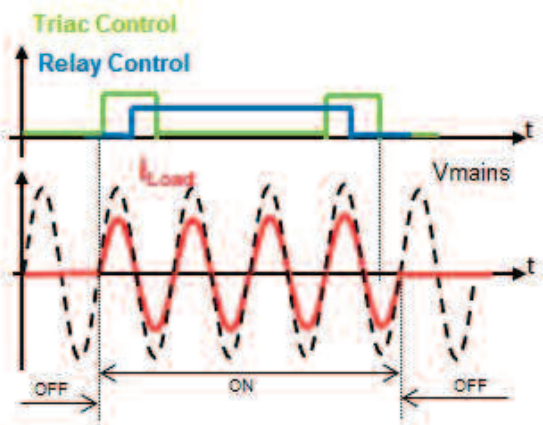

Figure 1: (left) motor-starter with hybrid-relays, and (right) relay/Triac control sequence

Figure 2: (left) Opto-Triac driving circuit, and (right) voltage spikes at current zero crossing

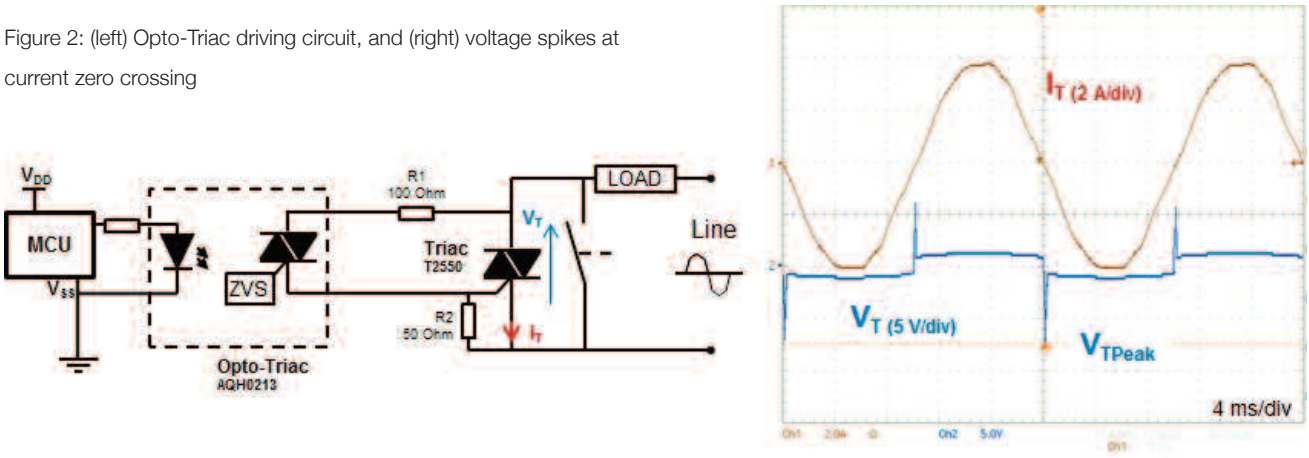

be placed in series with Line L1 in case the load neutral is connected.

Zero-voltage switching

Switching mechanical relays at a voltage close to zero can extend their life-time by a factor of ten. This factor would be even higher if switching occurs with DC current or voltage. Importantly, since the RoHS Directive (2002/95/EC) exemption for cadmium suppression expires on July 2016, silver-cadmium-oxide, which is used in contacts to prevent corrosion and contact welding, could be replaced with Ag-ZnO or Ag-SnO2. These contacts could present a shorter life-time unless bigger contacts are used to compensate.

Switch-on at zero voltage also allows the inrush current to be reduced with capacitive loads like electronic lamp ballasts, fluorescent tubes with compensation capacitor or inverters. This will also help to extend capacitor life-times and to avoid mains voltage fluctuations.

Additionally, solid-state technology allows the implementation of a progressive soft-start or soft-stop. A smooth motor acceleration and deceleration will reduce mechanical system wear and avoid damage to applications like pumps, fans, tools and compressors. For example, water hammer phenomena will disappear in pipe systems, and V belt slippage could be avoided, as could jitter with conveyors. Such HR starters are commonly employed in applications in the range of 4 to 15kW, but could also be used in applications up to 250kW.

Hybrid Relays are also used in heater applications. Heating power or room/water temperature is usually set with a burst control. A burst or cycle-skipping control consists of keeping the load on during a few cycles, 'N' and off during 'K' cycles. The 'N/K' ratio defines the heating power like the duty-cycle does in Pulse Width Modulation control. The control frequency here is lower than 25-30Hz. But this is fast enough for a heating system's time constant.

Source of EMI noise

Different control circuits could be considered to drive triacs, but an insulated circuit is mandatory in this application. As Figure 1 shows, the triacs do not have the same voltage reference, which is why an insulated control circuit — typically implemented with an opto-triac or a pulse transformer — is normal.

Figure 2 illustrates an opto-triac driving circuit. Triac gate current is applied through R1 when the opto-triac LED is activated (when the MCU I/O pin is set at high side). Resistor R2, connected between triac G and A1 terminals, is used to derivate the current coming from the opto-triac parasitic capacitor each time a voltage transient is applied. Usually a 50 to 100Ω resistor is used. The operation principle of this circuit makes a spike voltage occur at each Zero Current Crossing point (as shown on Figure 2), and this happens even if an opto-triac with a built-in Zero-Voltage-Crossing circuit is used.

Indeed, with an opto-triac circuit, a voltage has to be present across triac A1 and A2 terminals to

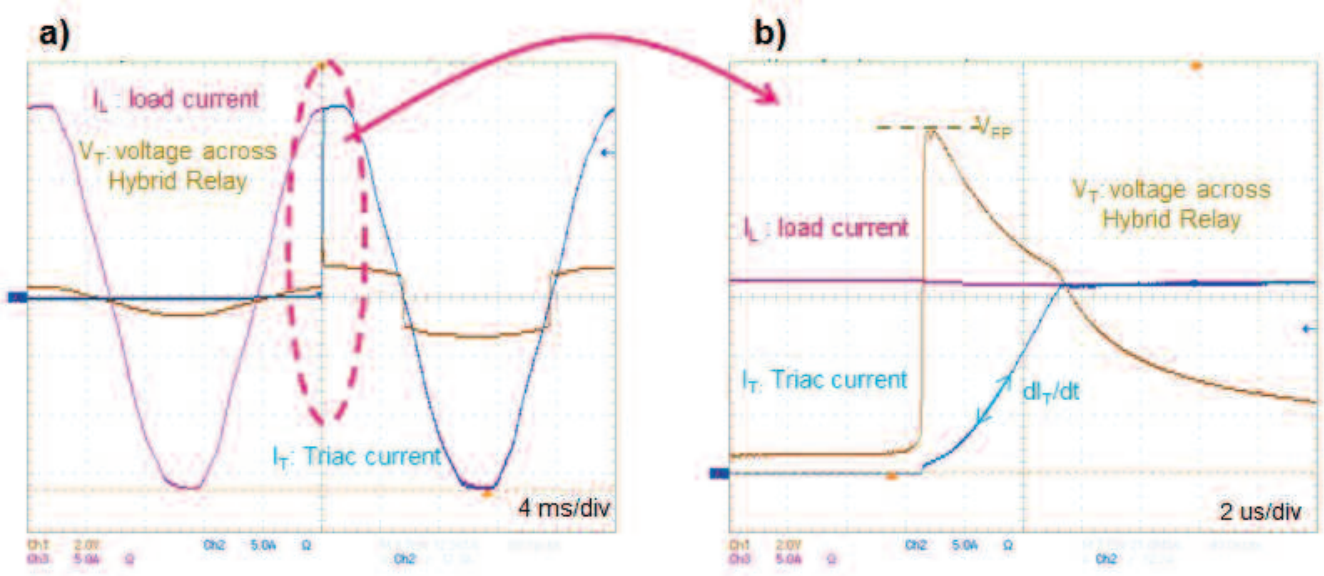

Figure 3: HR turn-off (a) – zoom at Triac turn-on (b)

allow a gate current to be applied. When the triac is ON, the voltage drop across it is close to 1 or 1.5V. This voltage drop is not high enough to provide a current through the gate as it is lower than the sum of the opto-triac and G-A1 junction voltage drops (both higher than 1V). So, each time the load current reaches zero, since no current is applied to its gate, the triac turns off.

Since the triac is off, the line voltage is applied back to its terminals. This voltage must then reach V_{TPeak} or high enough so that the current applied to the gate reaches the triac IGT. With a T2550-12G triac (a 25A 1200V triac with a 50mA IGT) used during the tests shown in Figure 2, this max peak voltage equals 7.5V (during negative transitions). Assuming a 0.8V and 1.1V typical voltage drop respectively for the G-A1 junction and for the opto-triac, this gives a 28mA gate current with a 200Ω R1 resistor. This current is the required IGT current for a turn-on in quadrant 3 (negative V_T voltage and negative gate current) for the sample used. The VTPeak voltage could even be higher for a sample with an IGT level closer to the max specified value (50mA) or if the device works at a lower junction temperature, as the IGT increases when the temperature decreases.

Since this V_{TPeak} voltage occurs at twice the line frequency (100Hz for a 50Hz mains), the emitted EMI noise could make the application exceed the emission limits defined by the EN 55014-1 standard for household appliances and electric tools. It should also be noted that this noise only occurs when the triac is conducting. As soon as it is bypassed by the relay, this noise disappears. The EN 55014-1 limit to apply for a discontinuous disturbance, depends on the repetition (or 'click') rate, i.e. on the HR operating frequency, and on the disturbance duration.

To avoid these voltage spikes, a pulse-transformer could be used in preference to an opto-triac. Adding a rectifier full-bridge and a capacitor to smooth the rectified voltage at the transformer secondary allows a DC current to drive the triac gate. Thus, there is no more spike voltage at each Zero Current crossing. However, a disturbance still occurs at the conduction transition from the mechanical relay to the triac. Such a transition only occurs at HR turn-off. Figure 3a shows the voltage spike that occurs during this phase; it happens precisely when the triac is switched ON — when the entire load current suddenly switches from the relay to the triac. Figure 3b illustrates a zoom-in view of the current increase through the triac. The dIT/dt rate is close to 8A/µs. As the Triac was triggered but not conducting (as the entire current was still circulating through the mechanical relay), its silicon substrate presents high resistivity when the current begins to flow. This high resistivity will lead to a high peak voltage, which equals 11.6V on the experimental test performed with a T2550-12G shown in Figure 3.

After the triac has started to conduct, both top and bottom P-N junctions of the triac silicon structure will inject minority carriers into the substrate. This injection will allow the substrate resistivity to decrease, and the on-state voltage to decrease down to approximately 1to 1.5V.

This is the same phenomenon that leads to a peak voltage drop across a PIN diode that turns on with a high rate of current increase. This is the reason why a PIN diode datasheet gives a V_{FP} peak voltage, depending on the applied dI/dt, which can have an impact on an application's efficiency if it occurs at a high frequency. For an HR application, this V_{FP} voltage only occurs once at HR turn-off and does not have to be considered to evaluate the power losses.

It should also be noted that, since the V_{FP} phenomenon is due to the time needed to modulate the substrate resistivity by injecting minor carriers, this voltage is higher for a 1200V device than for an 800V triac, like a T2550-8 for example. So the voltage a device is required to withstand has to be selected with care as an excessively high margin will lead to a higher peak voltage at turn-on.

Even if the measured peak voltage is higher than the one measured with the opto-triac circuit, EMI content is reduced as this phenomenon occurs only once per cycle, at each HR turn-off, and lasts just a few microseconds. Because of this, a pulse-transformer driving circuit is preferred despite its bigger size and its higher cost due to expensive ferrite cores.

Reducing peak voltage

To reduce the V_{FP} phenomenon on HR applications, a few easy tips can be implemented on the control circuit. The most effective one is to control the relay to switch OFF during negative current conduction. Indeed, the V_{FP} phenomenon is lower for a negative current. Figure 4 shows the V_{FP} voltage measured for the same test conditions as that of Figure 3b but for a negative current. It can be seen that the V_{FP} is reduced by 2, from 11.6V for a positive current to 5.5V here. The lower V_{FP} voltage for negative polarity is due to the easier silicon structure turn-on in quadrant 3 compared to quadrant 2 (positive A2-A1 voltage and negative gate current).

A second tip consists of increasing the triac gate current. For example, with a T2550-12G triac, the V_{FP} can be reduced by 2 to 3, especially for a positive switched current, when a 100mA gate current is applied instead of applying the specified IGT level (50mA) only. Another solution to reduce V_{FP} voltage is to try to open the relay close to the zero current crossing point. Indeed, limiting the switched current will also limit the applied dIT/dt at triac turn-on. Of course, to implement such a solution, a mechanical relay with a turn-off time lower than a few ms has to be selected. Reducing the dIT/dt can also be achieved by adding an inductor in series with the triac. A short PCB track between the mechanical relay and the triac is then not advised here.

More often now, Hybrid-Relays are used in applications to deliver longer life-times while achieving more compact size, a point particularly required in switching gears. Reasons for the occurrence of voltage spikes have been explained. Solutions have been presented to reduce their amplitudes like switching-off the relay for a negative current conduction, applying a DC and higher gate current, or adding an inductor in series with the triac.

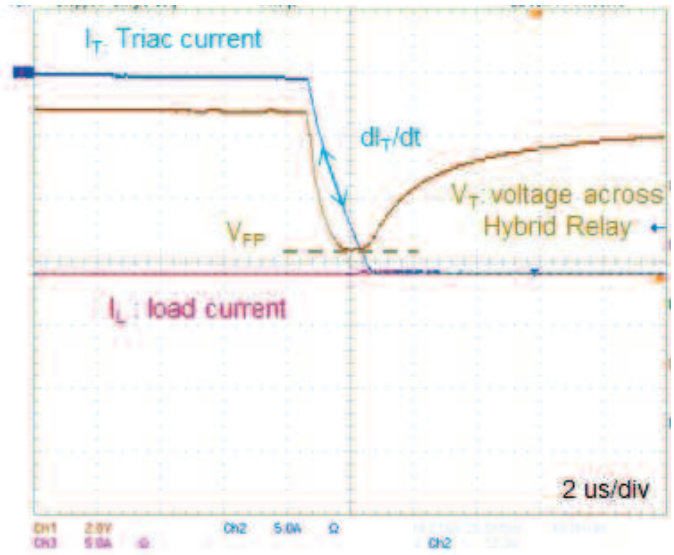

Figure 4: V_{FP} for a negative switched current

Challenges of Implementing Wireless Charging in Automobile Design

Keeping our handsets and other portable devices fully charged, on the go. By Peter Riendeau, Marketing Communications Manager, Melexis

As the sophistication of portable consumer electronics goods continues to increase, with higher degrees of functionality and more expansive feature sets, the rate at which they drain their batteries is (once again) becoming a limiting factor. Early mobile phones, back in the mid-Nineties had battery lives of just a few hours, but as the industry sector matured mobile handsets and other items of portable gadgetry were introduced that had the ability to last for days on end without a recharge. The move to 3G and 4G mobile communications has meant that the scope of what we can do with our handsets has changed dramatically, allowing us to indulge in online gaming, watch video content, access social media and more. A recent study by Ofcom found that younger adults (18-24 years old) were on average spending nearly 2 hours a day on their smartphone. This change in our lifestyle has meant that the time between handheld device recharges has shortened significantly. As a result there is a clear need for a means by which charge levels can be topped up during the day.

The advent of wireless charging has the potential to revolutionise the way handheld electronic devices are used on a daily basis. It will enhance the user experience and lead to greater convenience, as the devices' power budgets will be a lot less restrictive and there will be no need for users to carry different proprietary cables around for each item of electronic equipment. Wireless charging in homes, offices and public spaces certainly looks to be an attractive prospect, but there is one particular setting where it is likely to prove highly advantageous - in cars.

In-vehicle wireless charging has been discussed for several years and is now starting to see a major acceleration in its deployment. Though it follows the same basic set up as wireless charging infrastructure in other types of location, there are certain nuances that need to be taken into account that will distinguish it from normal implementations.

A powerful argument

Common wireless charging standards have built-in handshaking functions via which they can detect compatible handheld devices and then initiate charging. The sequence of polling so introduced will draw a relatively large amount of energy, which drains the car battery and may leave the car unable to start after being parked for a few weeks. At the same time, keeping RF emissions inside the vehicle at as low a level as possible is another concern, especially given the increasing number of wireless sources present there, with Wi-Fi hotspots, mobile phone signals and such like to contend with, as well as the electromagnetic interference intrinsic to an automotive environment.

Near Field Communication (NFC) may hold the key to its effective implementation in automobiles. This ultra-low power, secured, short range wireless communication standard is being deployed in a broad spectrum of applications, including point of purchase and access control. NFC is on its way to becoming almost ubiquitous in the handheld device market, with nearly all of the popular smartphones already employing it or about to

have it designed-in, including the Nokia Lumia, Samsung Galaxy, LG Nexus and most recently the iPhone 6. Likewise NFC is set to see widespread proliferation inside the car models themselves. Here it will perform a number of roles, such as automatic identification of the user and pairing of the mobile device to the vehicle's media centre (e.g. Wi-Fi or Bluetooth technologies). The NFC technology can also be used for access and start control, allowing NFC enabled handsets to be the 'key' to secure entry and drive-away of an automobile. This is just the beginning though, as there is also the prospect for it furnishing car manufacturers with a method by which to enable greater vehicle personalisation. If the vehicle occupant's mobile device is docked into the dashboard, then saved preference data can be used. Once automatic identification of the user has been completed the download of personal settings can be undertaken. This can allow the interior lighting or the seating positioning to be adjusted automatically to suit their liking.

With NFC being employed for these tasks, it makes good sense for it to also serve as the route through which wireless charging is also initiated. It offers a very low power RF link that can be kept running all the time within the cabin. Through this the vehicle can detect if a smartphone or other handheld device is present and then carry out handshaking and begin the wireless charging process.

Through the combination of the respective expertise in automotive grade NFC transceiver technology and wireless charging, Melexis and Freescale Semiconductor have introduced a jointly-developed reference design (Figure 1). It includes all the hardware and associated software drivers necessary to permit fully autonomous operation of the wireless charging subsystem. The embedded NFC stack is controlled via a standard compliant NCI (NFC Controller Interface) and driven by a standard compliant upper NFC software stack (e.g. NFCStack+ from Stollmann). It deals with the coexistence of both Wireless Charging Transmitter (WCT) and NFC technologies on a single and easy-to-use module.

Getting a head start

Melexis has put itself at the forefront of NFC roll-out within the automotive arena with its MLX90132 multi-protocol transceiver IC, which is a fully automotive compliant (AEC-Q100 qualified) device that operates at 13.56MHz. The NFC analog part embeds a minimum set of components around the MLX90132 to make it functional as an NFC device, supporting the following communication modes: Reader/Writer, Peer-To-Peer (Initiator and Target) and TAG Emulation. It allows the transmitter to provide up to 300mW RF power to an appropriate antenna load. In the reference design this is complemented by a

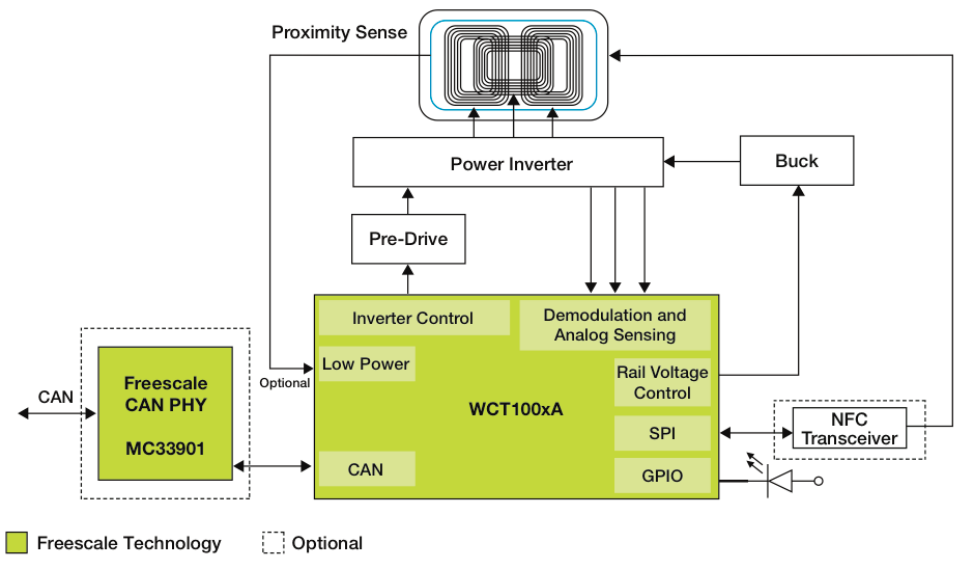

Figure.1: The WCT-5WTXAUTO NFC/Wireless Charging Reference Design

MWCT1003AVLH microcontroller from Freescale.

The microcontroller supports the NFC low-layer stack and incorporates Freescale's own unique wireless charging mechanism. The issue of co-existence is dealt with in hardware thanks to a specially designed on-PCB printed NFC antenna with an optimised matching network taking into consideration the detuning that results from the presence of wireless charging coils. The proposed architecture, including both stacks in one common microcontroller, allows the implementation of dedicated software handshaking to reinforce the co-existence of the two technologies. A first software version allows the sequential operation of the NFC and wireless charging activities, subsequent releases will allow the possibility to further enable transparent and simultaneous operations.

The reference design's embedded field and tag detectors mean that it has extremely low power consumption, reducing to a minimum the power dissipation of the final application. The Electro-Magnetic Disturbance (EMD) algorithm improves performances in both Reader/Writer and Peer-to-Peer modes.

Automotive manufacturers are looking to integrate NFC technology into their new models in order to take care of a multitude of different tasks. With it being an integral part of both handsets and vehicles it presents automotive engineers with the means by which a variety of functions can be addressed, including wireless charging initiation, using an extremely low power wireless communication standard instead of having to rely on a more power hungry alternative.

Reader Offer!

Win a Microstick for 3V PIC24F K-series

Electronic Specifier Design is offering its readers the chance to win a Microstick for 3V PIC24F K-series (#DM240013-1). This Microstick is an easy-to-use, flexible USB powered development platform. It's the perfect solution for those looking to get started with Microchip's lowest cost 16-bit microcontroller families, PIC24F "KL" and "KA", designed for extremely cost sensitive consumer, medical and industrial applications.

The Microstick for 3V PIC24 K Series is designed to provide an easy-to use, economic development environment for 16-bit microcontrollers. The board includes an integrated debugger/programmer, a socket for the target microcontroller chip and pins that facilitate insertion into a prototyping board for extremely flexible development.

At about half the size of a credit card, it is extremely portable, and can be plugged into a prototyping board. This Microstick ships with a USB cable, header pins for proto board use, and the PIC24F16KL402, and PIC24F16KA102 MCUs.

For the chance to win a Microstick for 3V PIC24F K-series, please visit http://www.microchip-comps.com/esd-3v24fk and enter your details in the online entry form.

C what you get

In August this year, the USB-IF released the specification for the Type C connector for USB 3.1, promising a simpler and more integrated single-cable solution. Steve Rogerson looks at how likely that might be

Offering 10Gbit/s and up to 100W, USB can now be used not just for powering smartphones and the like but for desktop PCs and workstations. The Type C connector is smaller, making it more suitable for thinner devices, and more consumer-friendly as it can be inserted any way up. A Type C cable will also have the same connector at each end so it can be plugged in either way. And there are more pins, which means the USB connection can be used for extra applications.

"The real focus for USB is to deliver a single source for audio, video, data and power all over a single cable and connector," said Jeff Ravencraft, President of the USB-IF. "It is bidirectional, it doesn't matter which way you plug it in and which end you plug in. There is no longer an A and a B. With the Type C connector we have a very robust connection." The system is backwards compatible but only through the use of dongles to bridge between legacy USB connections and Type C cables, but Ravencraft envisages that eventually the C connector will take over everywhere.

And Gordon Lunn, Global Customer Engineering Support Manager for FTDI Chip, added: "The new Type C connector signalling is fully backward compatible, so engineers looking to make use of it won't have to rework there designs, they will be able to make use of the same ICs and firmware that they were previously. A full range of conversion cables will be available (B-to-C, A-to-C, etcetera) so implementation of this new interconnect option will be simple enough to accomplish." But Alan Jermyn, Vice President for European Marketing at Avnet Abacus, said he had been speaking with suppliers and the feedback he was getting suggested they were worried about backwards compatibility issues.

"Our engineering team has been monitoring the USB-IF," he said. "The connector is reversible so it can be plugged in either direction similar to Apple's Lightning connector. It is not backwards compatible with any previous USB connectors, which is why they have specified what looks like nine new cable types each with a different legacy connector attached. The lack of backwards compatibility is a concern. Realistically, they are up to two years out from offering these in production."

One of the traditional ways to connect multiple devices via USB was to use a hub plugged into the mains. This is now moving to having the hub built into the display, which is connected to the mains. This can then be used not just to connect the other devices but to provide power to them as well. The 100W power delivery is a significant increase over previous versions. USB 2.0, for example, could deliver 500mA and 5V, or 2.5W. With 3.0, this went up to 900mA, or 4.5W. The added battery charging specification boosted this to 7.5W. Data delivery is also twice as fast as USB 3.0's 5Gbit/s and five times more than Hi-Speed USB. Also, the voltage, current and

Gordon Lunn: "Implementation of this new interconnect option will be simple enough to accomplish."

Jeff Ravencraft: "Maybe it could have been a little smaller but I think it is the right size."

direction of the power flow can be negotiated.

"The direction can be changed in real time," said Ravencraft. "For example, a netbook would normally be a provider of power and the cell phone connected a consumer of power. But connect the netbook to the display and that could provide power to the netbook. The flow of power can be changed dynamically." Using the latest power delivery method with the Type C connector also brings extensions beyond power. The most noticeable for most users will be the dual role data transfer. Like the reversible power, the status of host and device can also change dynamically. For example, if a phone and notebook are connected, the phone would be a device, but if you connect the phone to, say, a TV or printer it could become the host.

Alternate modes

The other key change is discovering and configuring what are known as alternate modes. These let vendors run additional features on the connectors, such as PCI Express for example. "The USB-IF is developing guidelines so you can easily recognise the capabilities you are getting," said Ravencraft.

How this works is that the 3.1 data channels are on one side of the 24-pin connector, and that can be either side depending on which way the connector is inserted. The other side of the connector also has high-speed data pins and it is these that can be used for other purposes.

"You could use them to transmit video using embedded DisplayPort," said Tom Watzka, Technical Marketing Manager at Lattice Semiconductor. "There is a protocol that lets you set this up between the host and the client, but it is a challenge to enable this capability. The specification is very new and there are no solutions that we are aware of. The opportunity is that these specifications are really flexible. This is what is driving the interest. A lot of OEMs see the possibility of using this for their specific needs."

Not surprisingly, he sees an FPGA route to helping companies handle this flexibility. "With our FPGAs, we have begun development on this," he said. "Because they are programmable, we have been able to adapt them very quickly. There is a lot of flexibility with the kind of data and that can be used for competitive advantage."

Tim Wang, Product Marketing Manager for Lattice, also sees FPGAs as a way to handle the other features of Type C: "The FPGA can do all the work," he said. "It can figure out which is the host and which is the slave. You need to figure out what kind of power the slave requires, and the FPGA can do that. And it can do the negotiations with the host."

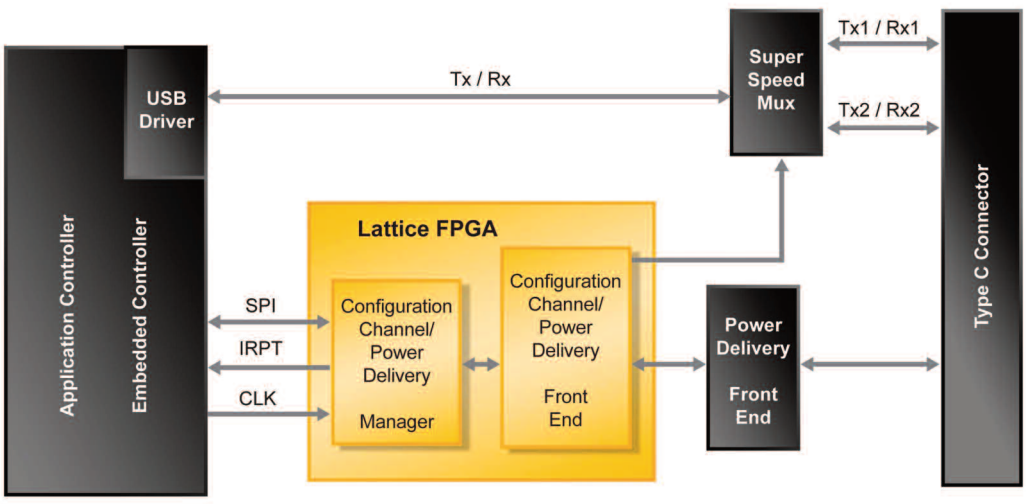

How an FPGA can be used to handle the flexibility with the Type C connector

Drivers and dissent

So what were the drivers for these changes? First, the

Type A connector was too large for emerging computer platforms; it needed to be thinner. The Micro-AB connector might have been good enough for mobile phones, but it didn't have the robustness needed for other roles. Also, many consumers were confused with plug orientation and cable direction – trying to force a USB connector into a socket the wrong way up is not uncommon. The Type C connector addresses all these issues.

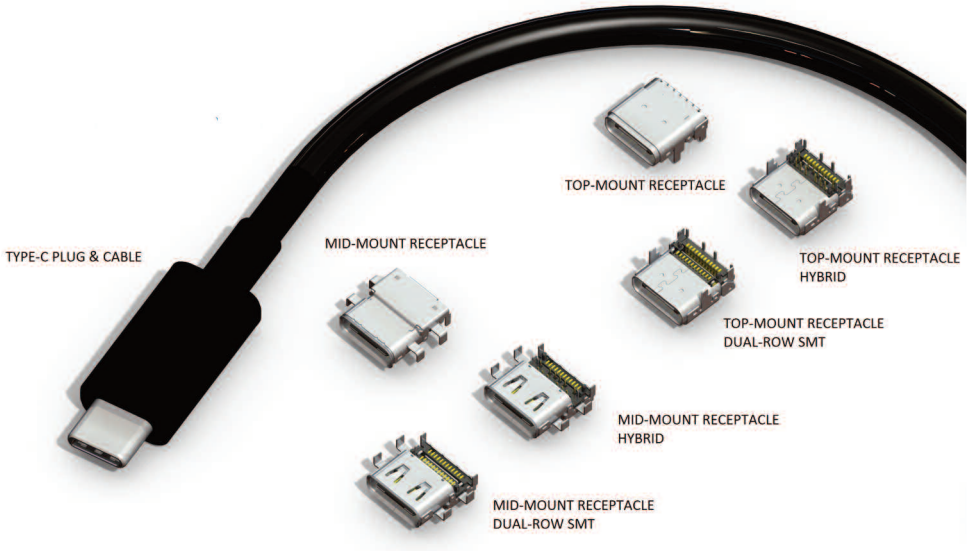

USB Type C connector with receptacles (photo from Foxconn)

The specification was finalised in August and was the result of an industry-wide collaboration. "Anybody could participate and help define it," said Ravencraft. "We are now compiling a compliance and certification programme and we want to have something in place for early 2015. We expect to see products in the market in early 2015."

However, not everybody is happy. Molex, which already has samples of its Type C connector and is planning product in 2015, is a little worried on the size and cost of the connector. "Initial feedback from the mobile phone makers is that the connector size of Type C is too big," said Carol Liang, Group Product Manager for Molex in Shanghai. "Connector width and heights are industry standard dimensions but we have generated different proposals to shorten connector depth (24 pins with dual row SMT soldertails or 14 pins with single row of SMT soldertails). The dual row SMT soldertail design, however, may pose inspection and rework concern to mobile phone makers."

Ravencraft responded: "There was a lot of work put into the size. There had to be give and take. Some wanted it smaller, but we had to ensure the robustness, both mechanical and in terms of shielding for RFI and so on. The industry has agreed on this size. Maybe it could have been a little smaller but I think it is the right size." Liang also has concerns about cost. "Early indications are that the total solution may cost more than Micro USB 2.0," she said. Ravencraft responded: "Initially, with new technology there will be a slight price difference. But even from the outset, they will be competitive. But for the long term, with the type of volumes for USB the price will become very inexpensive."

However, Liang thinks this will only happen with EU support, which she believes is not certain. "The EU has adopted Micro USB 2.0 as the standard IO for charging mobile equipment and this makes the demand on Micro USB 2.0 very high," she said. "At this early stage the industry is unclear as to whether the EU is going to do the same thing with Type C. If the EU does, we would expect volumes to be high. However, cost pressure will also be high. Hence, we need to design Type C with both low and high volume manufacturing strategy planned out now so that we can be cost competitive down the road should volume demand pick up."

On the face of it, the Type C connector looks just what the industry has been looking for. It is smaller, easier to use and brings the flexibility for other high data applications. But there is worry that dongles and special cables will be needed to handle the large numbers of equipment with existing USB sockets. And some believe that the size is still too big for modern devices. And, yes, it will cost more and that cost will only come down if the other problems fail to stop widespread adoption.

Simplifying Standards

Supporting multiple high-speed transmission interface standards and accelerating time-to-market with economical SERDES enhanced FPGAs. By Benny Ma, Applications Engineer, Lattice Semiconductor

The convergence of field programmable gate array (FPGA) and high-speed SERialiser/DESerialser (SERDES) technologies has led to the emergence of SERDES-enhanced FPGAs as a cost-effective alternative to ASICs in applications which require a multi-Gigabit data link across a printed circuit board, backplane, or cable. This new class of programmable devices is changing the way many products are designed, particularly as these devices are becoming available at increasingly lower-power and lower-cost.

Originally developed for high-performance packet processing and routing elements in carrier-class networking equipment, FPGAs equipped with embedded SERDES transceivers have been available for over a decade. Combining low-power multi-Gigabit SERDES transceiver cores with a price-optimised FPGA architecture creates a versatile solution platform for applications which make extensive use of SERDES-based interfaces such as Ethernet (XAUI, GbE, SGMII), PCIe (PCI Express) SRIO (Serial RapidIO) and Common Public Radio Interface (CPRI). These include small-cell wireless infrastructure, microservers, broadband access, industrial video and other high-volume applications where low cost, low-power, and small form-factors are key design constraints.

Many of the early applications to embrace SERDES-enhanced FPGAs have been innovative RF, baseband and backhaul products which address the needs of wireless network operators. Virtually every carrier throughout the world is making massive investments to upgrade their infrastructure as they struggle to keep pace with the explosive growth in demand for mobile data and video traffic. In addition to upgrading their existing base stations to support high-capacity 3GPP/4G wireless standards, many operators are embracing a new Heterogeneous Network

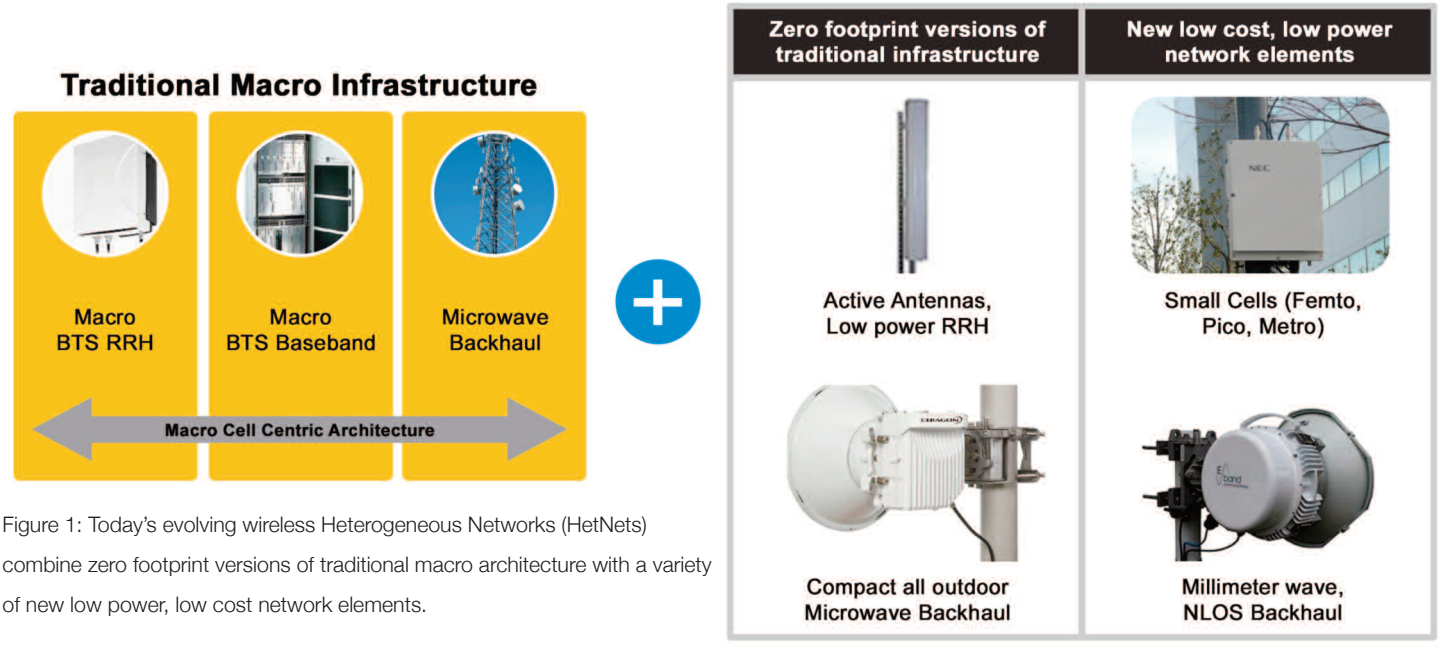
Traditional Macro Infrastructure
Macro BTS RRH
Macro BTS Baseband
Microwave Backhaul
Macro Cell Centric Architecture
Zero footprint versions of traditional infrastructure
New low cost, low power network elements
Active Antennas, Low power RRH
Small Cells (Femto, Pico, Metro)
Compact all outdoor Microwave Backhaul
Millimeter wave, NLOS Backhaul

Figure 1: Today's evolving wireless Heterogeneous Networks (HetNets) combine zero footprint versions of traditional macro architecture with a variety of new low power, low cost network elements.

Messe München International

Connecting Global Competence

Welcome to Planet e.

The entire embedded universe at a single location!

Tickets & Registration
www.electronica.de/en/tickets

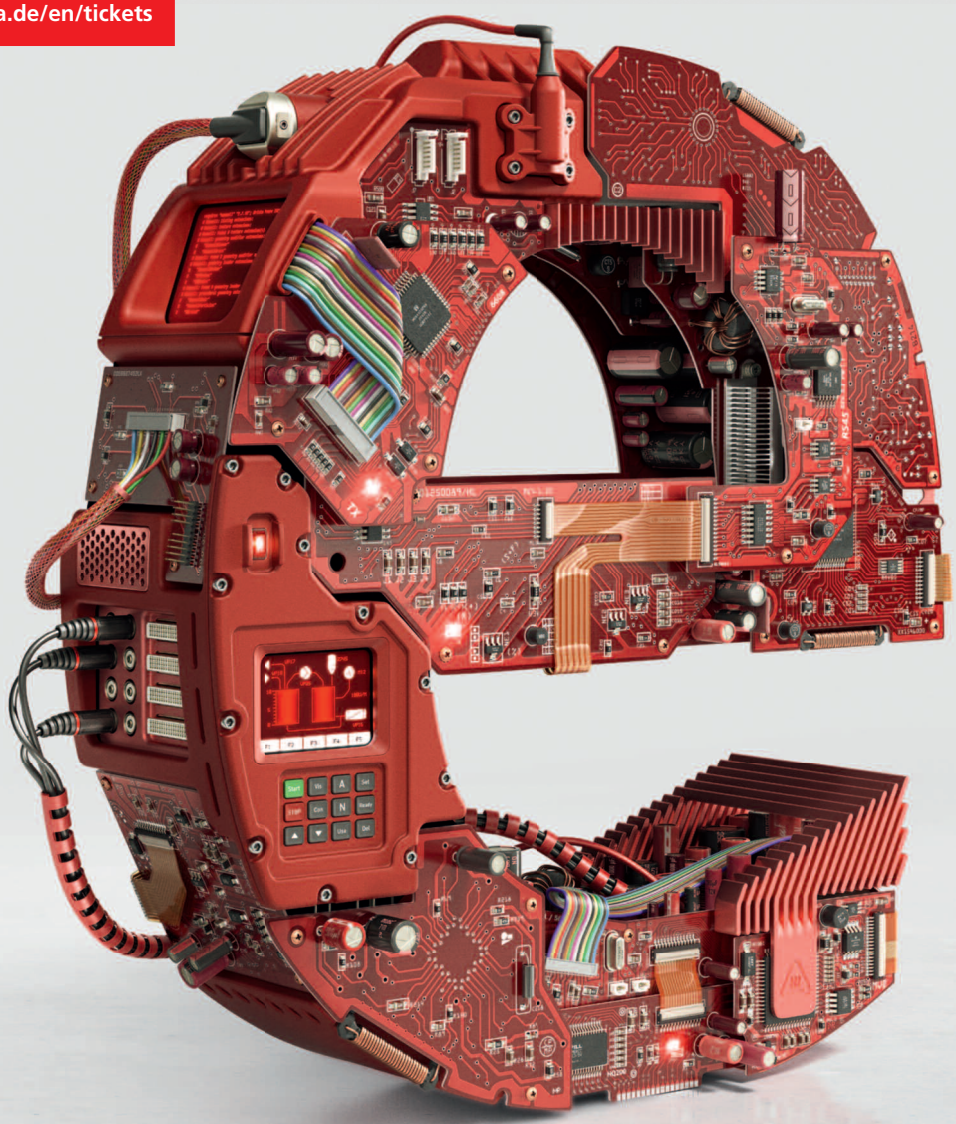

26th International Trade Fair for
Electronic Components, Systems
and Applications
Messe München
November 11–14, 2014
www.electronica.de

50 years
electronica

electronica 2014
inside tomorrow

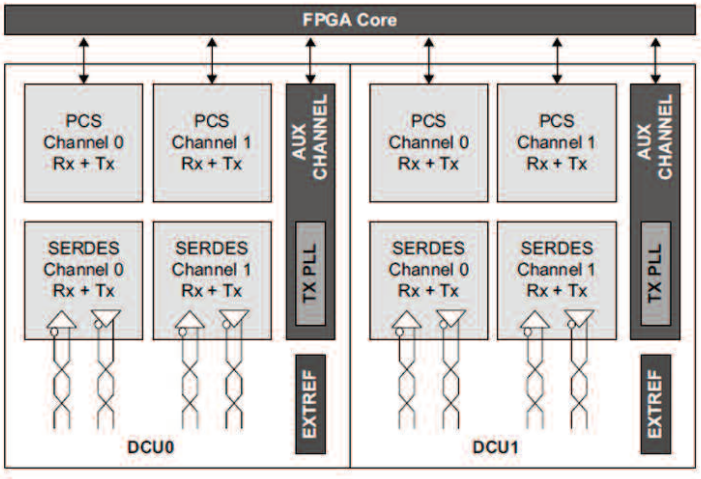

Figure 2: Block diagram of the dual-channel SERDES element used in the ECP5 FPGA series.

(HetNet) architecture in which the traditional macro infrastructure is supplemented by a new class of low power nodes (LPNs) such as small cells, low power remote radio heads, and relays (Figure 1).

These compact, low-power (typically between 100mW and 10W) wireless nodes can add capacity in high-traffic areas or extend wireless coverage into buildings, public spaces, and urban canyons which are beyond the reach of conventional base stations. This requires LPNs to be highly configurable in order to support multiple air interface standards and RF frequencies, compact and rugged enough to blend into the urban landscape, and inexpensive enough to justify widespread deployment.

In any LPN, FPGAs with SERDES capabilities can be used to implement data path bridging and interfacing and the packet-based network interfaces (GbE, 10GbE/XAUI, XGMII) commonly used to connect small-cell clusters with the backhaul infrastructure. They can also be used to implement the XGMII interface and most of the digital functionality in smart SFP (small form-factor pluggable) transceiver products, commonly used in broadband access equipment.

Like most SERDES-Enhanced FPGAs on the market today, Lattice Semiconductor's ECP5 FPGA family uses embedded SERDES transceivers which provide the baseline functionality required to support most commonly-used high-speed serial interfaces. The ECP5 series includes devices with varying amounts of programmable logic elements (up to 85k LUTs) and up to four transceivers. The transceivers are implemented as pairs within a Dual Channel Unit (DCU) containing a single shared reference clock input and Tx PLL. Each SERDES channel can be configured independently to perform high-speed, full-duplex serial data transfers at data rates from 270Mbit/s to 3.2Gbit/s (Figure 2).

Data coding

The Physical Coding Sublayer (PCS) logic in each channel can be configured to support several types of Ethernet interfaces and several other common networking and system interconnect standards. Since transceiver power consumption varies according to how it's configured and which features are used, no single number could be considered accurate. But, as a first-order approximation, it's safe to say that in most simple configurations, a single channel 3.25Gbit/s SERDES consumes less than 0.25W. Quad-channel SERDES elements which support similar functions often consume under 0.5W. Different combinations of protocols within a DCU are permitted subject to certain conditions as specified in the ECP5 SERDES Usage Guide.

Soft IP can be used in conjunction with the SERDES channels to support protocol level function for high-speed serial data links such as PCIe, CPRI, SD-SDI, HD-SDI and 3G-SDI. For custom applications, the user can use his own protocol level logic, giving users full flexibility to design their own high-speed data interface. The PCS also provides bypass modes that allow a direct 8-bit or 10-bit interface from the SERDES to the FPGA logic, allowing user to implement their own data coding.

In order to compensate for the inter-symbol interference (ISI), attenuation, reflection, and other phenomena which SERDES signals encounter as they traverse printed circuit

boards or cables, the transceiver uses a combination of transmitter and receiver equalisation techniques which can be programmed via the device's configuration registers. The transmitter's high speed line driver has adjustable amplitude settings and termination resistance values which can be optimised for attenuation due to the length of the channel, and adjusting termination resistance to match the channel trace impedance to minimise signal reflection. It can also perform transmit equalisation using adjustable pre/post-cursor de-emphasis settings which reduces inter-symbol interference (ISI) caused by interactions between the bit being transmitted and the energy from the previously-transmitted bit still present in the transmission line.

The receiver includes a linear equaliser (LEQ) which is used to selectively amplify the frequency components in the data rate range which tend to be more heavily attenuated over long runs across a PCB or backplane. Compensating for this frequency-dependent attenuation helps mitigate the inter-symbol Interference (ISI) which would otherwise occur in the receive signal. The receiver offers four levels of equalisation which can be selected according to each channel's transmission characteristics.

Many of the challenges of using SERDES-enhanced FPGAs arise from the characteristics of the Ball Grid Array (BGA) packages frequently used in the space-constrained applications where many of these devices are used. A key challenge in adopting fine-pitch (0.8mm or less) BGA packages is the design of a route fanout pattern that maximises I/O utilisation while minimising fabrication cost. At these contact densities, it becomes difficult to route traces out from inner balls without additional PCB routing layers. In addition, high-speed signals, such as SERDES or DDR3/LPDDR3 signals, require even closer attention to maintaining a uniform gap for controlling trace impedance, and matching lengths between paired traces for groups of signals on a Source Synchronous bus, insuring provisions for a proper ground plane and isolation layers in the PC board and other issues which maintain signal integrity.

Fine pitch packages offer advantages and disadvantages alike. Finer pitch means that the trace and space limits will have to be adjusted down to match the BGA. Many times a design can get away with small traces underneath the BGA then fan out with a slightly larger trace width. The PCB fabrication facility will need to be aware of your design objectives and check for the smallest trace dimensions supported. Smaller traces take more time to inspect, check and align. Etching needs to be closely monitored when trace and space rules reach their lower limit. The combination of fanout traces, escape vias, and escape traces that allow routing out from under the BGA pin array to the perimeter of the device are collectively referred to as the BGA breakout (Figure 3). The fanout pattern will arrange the breakout via, layer, and stack-up to maximise the number of I/Os that can be routed.

Packaging

Lattice has created a package design for the caBGA554 and caBGA756 form factors in 0.8mm ball pitch with a new package ball

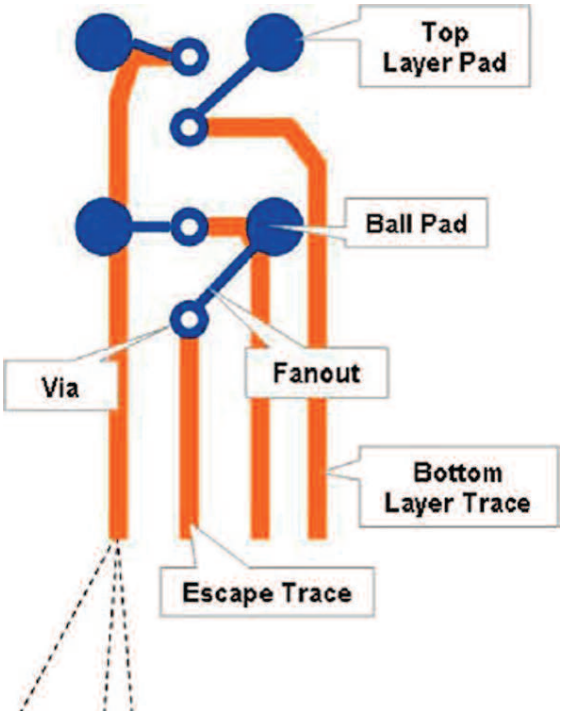

Figure 3: BGA breakout routing features.

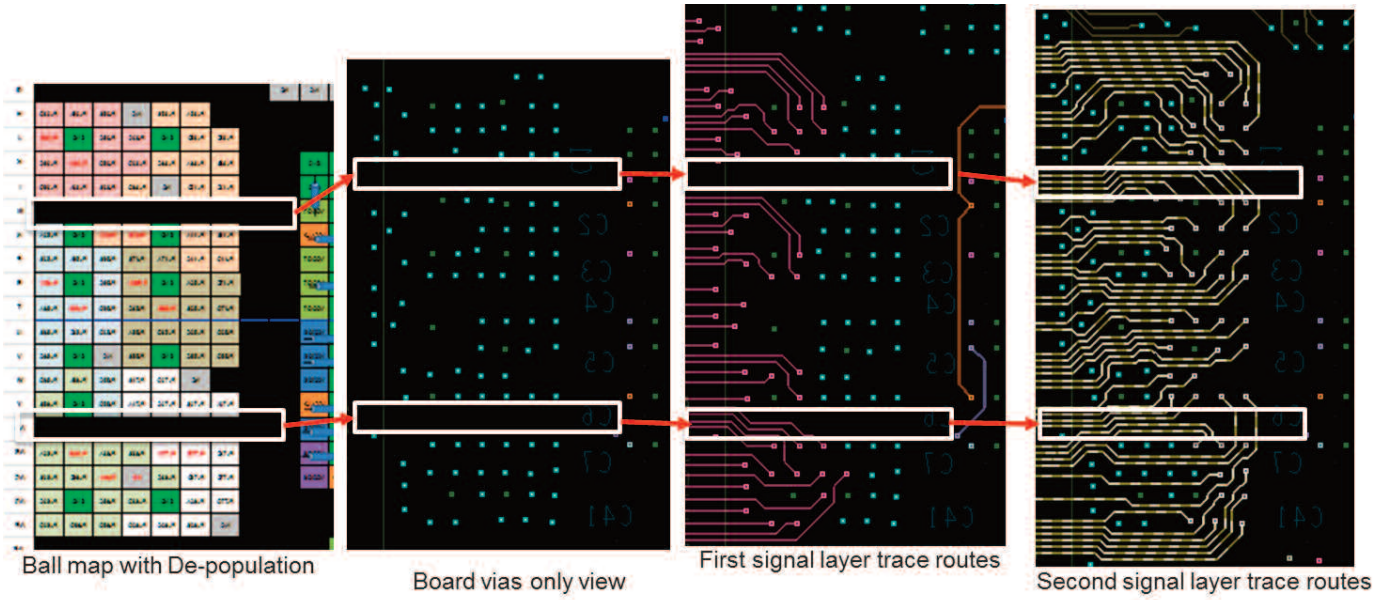

Figure 4: Top: Depopulation of BGA balls enables simpler, cleaner trace routing using fewer PCB layers. Bottom: Unobstructed routing of SERDES signal traces allows greater control over channel characteristics which affect signal integrity.

break-out scheme, which allows traces to be brought from inner row balls. The package design also selectively depopulates unneeded ball positions to open up real estate for easier routing. In addition, careful assignment of signal/power/ground balls provides better skew matching, lower cross-talk among busses of high-speed signals, as well as power/ground ball assignment that allows under the DUT low inductive decoupling capacitors for supply pins.

Although this generation of SERDES-Enhanced FPGAs has achieved significant reductions in power consumption, cooling can still be an issue, especially in products which are routinely deployed in outdoor environments. For this reason, the new packaging design also maximises the number of ground vias located in close proximity to the FPGA chip itself to provide maximum conductive heat dissipation to the PCB under device. The example board shown in the top image in Figure 4 illustrates a design which takes full advantage of the benefits of ball de-population while using only two signal PCB layers (even with thru-hole vias).

When a row of balls on the package is de-populated, it removes vias that are needed for those balls. This creates unobstructed area on both signal routing layers, which offers more flexibility on the signal routes to break out to the edge of the package. Of greater interest however, is the routing for the SERDES signals. The de-population of selected groups of balls enables unobstructed routing of the FPGA's high speed SERDES signals. Each pair of SERDES signal traces is closely matched in terms of length and maintains accurate trace-to-trace spacing to insure stringent control over impedance. Spacing for pair to pair, and pair to FPGA signals, is also closely controlled to minimise cross talk.

By combining programmable logic with high-speed serial data transceivers, SERDES enhanced FPGAs can support a wide range of networking and system interfaces while providing programmable logic elements which can supplement and in some cases eliminate ASICs and ASSPs used in conventional designs. Their programmable capabilities help enable rapid development cycles and make it possible to create easily-upgradeable products and configurable platforms which can support multiple networking and communication standards. As with any SERDES device, they also bring several challenges to the design process, mainly in the area of packaging, PCB layout, and signal integrity. But a combination of good design practices and new packaging technologies can be used to insure SERDES-Enhanced FPGAs are able to achieve their full potential.

Getting to 5G

The promise of seamless surfing comes with significant challenges. By Hervé Oudin, EMEA 5G Program and Cross-Division Wireless Marketing Manager, & Jan Whitacre, WW Wireless Program Lead, Keysight Technologies

The latest improvements to mobile networks have allowed new services and created a new set of end-user expectations, witnessed by the rise of more complex and data-hungry applications for smartphones and tablets. Some applications now support 'seamless connectivity', or the ability to continue using an application when moving between devices without interruption to the content. Providing this capability requires access to and control of the content over multiple networks; WiFi, cellular and broadband. With mobile data consumption currently forecast to almost double year-on-year for the next five years, the network operators maintain they will struggle to meet long-term demand without more spectrum.

To support the huge increase in the numbers of devices and performance requirements, studies describe the key network attributes that will be required: an integrated wireline/wireless network, where the wireless part is a dense network of small cells with cell data rates of the order of 10Gbyte/s made possible by high-order spatial multiplexing (MIMO). A round-trip latency of 1ms will give the ability to deliver the interactive high-resolution streaming video that's needed for 'immersive virtual reality' applications. It's now assumed that devices will support simultaneous use of multiple air interfaces, including not only extensions to today's RF cellular frequency bands, but also operation at microwave or millimetre frequencies. With these attributes, the combined network will support everything from simple M2M devices to next-generation phones, tablets and PCs, with the monitoring and control of literally billions of sensors and multiple simultaneous streaming services, while supporting the massive data collection and

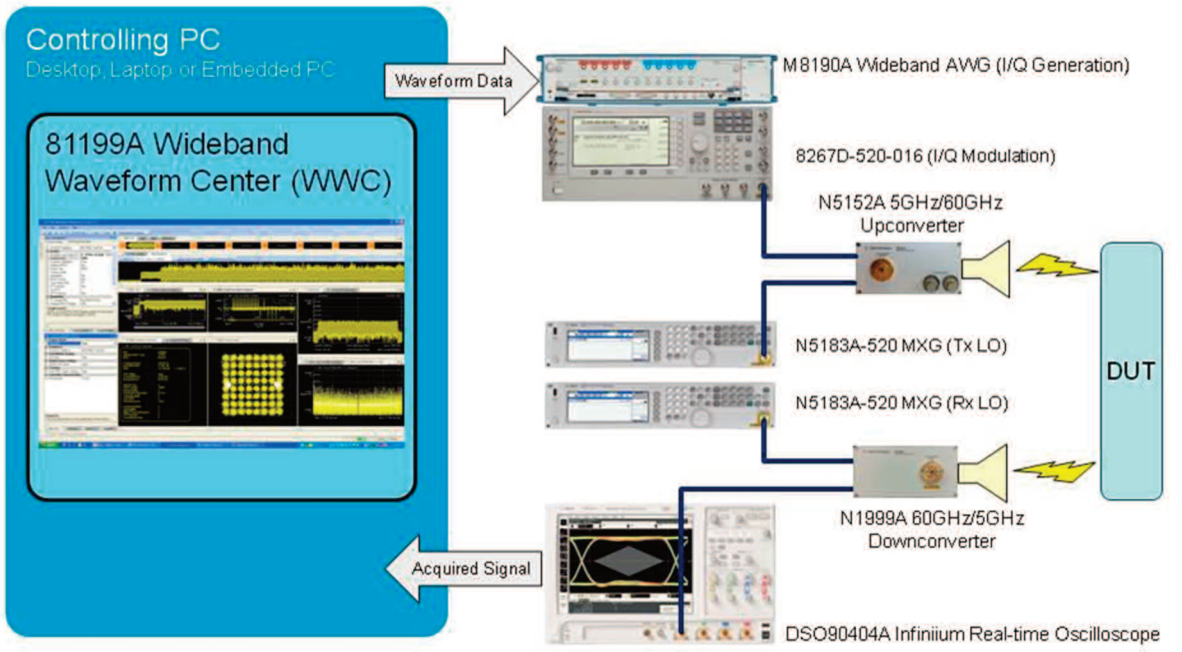

Figure 1: 60GHz stimulus/response test system with software.

distribution needs of the Internet of Things.

The 'everything everywhere and always connected' vision for 2020 and beyond that's presented in the studies for fifth generation networks assumes a number of new paradigms: devices can operate at frequencies from a few hundred megahertz to greater than 100GHz, indoor cell sizes that may be as small as a single room, and a dense network of pico- and femto-cells to maximise the number of users that can be supported. 5G's goal is to provide a high-capacity network capable of 10Gbit/s peak and always delivering 1Gbit/s rate to however many users want it; in other words to provide each user with 'infinite' bandwidth; all the bandwidth they need, anywhere, anytime including crowded areas such as sporting events and conventions.

None of the studies have specific details of the core network that joins everything together, but they all assume the seamless connectivity mentioned earlier is a given.

In the cellular world, capacity gains come essentially from three variables: more spectrum, better efficiency and better frequency re-use through progressively smaller cell sizes; all of these areas are being investigated. Researchers are looking at a wide range of frequencies from current cellular frequency bands to 28GHz, 38-40GHz, 57-64GHz, 70-75GHz, 81-89GHz and even 140GHz. They are looking at wide bandwidths from 0.5GHz to > 3GHz and antenna research with multiple-antenna configurations to increase capacity and focus cell performance in the direction of specific devices. Work continues on small cells and heterogeneous networks, with new techniques for self-organising networks, software defined radios capable of multiple air interface standards and software defined networks based on cloud computing are already being proposed for future 4G standards releases, and these will be extended to 5G. And they are investigating new physical layers such as GFDM, FBMC, UFMC, BFDM and NOMA.

Challenges

The World Radio-communication Conference, hosted by the International Telecommunication Union, is held every three to four years. Its mandate is to agree international radio frequency issues, including frequency allocation standards for mobile networks. The next WRC is scheduled to be held in Geneva in 2015 (ITU-R WRC-15) where initial discussions of 5G will be held.

Compared to previous generations of mobile network, 5G presents a number of new design and test challenges. Component and system design and test at microwave and millimetre frequencies has been around for many years, but its application to high-volume, low-cost devices for the consumer market is relatively recent. There's already an unlicensed frequency band being used at 60GHz for the wireless LAN standard 802.11ad, which features a 2GHz channel bandwidth. Similar use may be made of licensed spectrum in the 28GHz range, where Samsung and others have already reported experimental results, and in other ranges, where a number of university studies are under way.

Work is underway on channel modelling, to characterise how signals behave at mmWave frequencies. In 'real' user devices, these frequencies would likely use antenna components bonded directly to the transmitter and receiver chips, making connection to test equipment a challenge. This type of configuration has inherent issues in providing reliable and repeatable system calibration. Base station radios will typically feature antenna matrices for beam steering (directing RF attributes towards a specific device) and/or multiple transmit/receive streams (MIMO) for capacity enhancement. Testing user devices will mean emulating these real-world network conditions and test equipment suppliers will need to provide new channel measurements and simulation models for initial development and complex baseband and microwave sources for performance verification. Figure 1 shows a system designed for testing 802.11ad components at 60GHz and gives an idea of what might be needed for millimetre wave 5G design and development.

Any potential new physical layer technology (PHY) attributes are still to be decided, but it's likely that any 5G devices will need to operate in a number of different Radio Access Networks (RANs). The new

PHY will include new modulation and coding schemes that are more efficient at very high symbol and bit rates (e.g. the use of Golay sequences in 802.11ad). Challenges here include everything from battery power consumption (meeting user expectations) to supporting receiver systems that can demodulate and decode data from multiple carriers using different PHY characteristics simultaneously, then integrate the data into a single useful data stream.

Today, research into next-generation systems is being carried out in universities, either on their own or as part of consortiums or forums with support and direction of commercial partners, or in the R&D departments of network equipment manufacturers, chip and device manufacturers and network operators. While the standardisation process for next-generation may not start until 2015, it's expected that some of the research currently under way will be incorporated into the next-generation communications systems which will begin to roll out around 2020.

Blue sky thinking

With investigations being totally open to many areas, researchers need a wide range of solutions with much flexibility to cover all the frequency ranges and all the analysis needs. Today, Keysight provides full range of flexible simulation and measurement tools that bring insight to this research. Vector Network Analysers allow in-depth design and test of millimetre wave components up to 110GHz, such as the antenna array elements needed for beam-steering and MIMO. With anticipated devices from simple sensors designed for years of unattended use, to next-generation smartphones and tablets, battery life will be key to meeting user expectations. Keysight battery drain analysis systems offer designers a power management solution to test their devices under normal, high and low voltage conditions. For the physical layer, where advanced digital signal processing meets RF, SystemVue is a system-level design automation environment that accelerates design and verification of communications systems. It combines with Keysight measurement products to create an expandable environment for modelling, implementing, and validating next-generation communications systems. It enables a virtual system to be verified from the first day of a project, beginning with simulation models, and gradually incorporating more real-world measurements as the design is translated into working hardware. It can be used in conjunction with Keysight signal sources to create complex arbitrary waveforms to verify theoretical channel models once the design is implemented. SystemVue can also be used in conjunction with Keysight 89600 VSA software, a comprehensive set of tools for demodulation and vector signal analysis. VSA works with a range of measurement tools: benchtop and modular signal analysers, digital oscilloscopes and wideband digitisers, to match the frequency and bandwidth of the signal to be analysed. Together these measurement, simulation, and signal generation and analysis tools enable the exploration of virtually every facet of the components and signals that will become part of the advanced designs needed for next-generation communications systems.

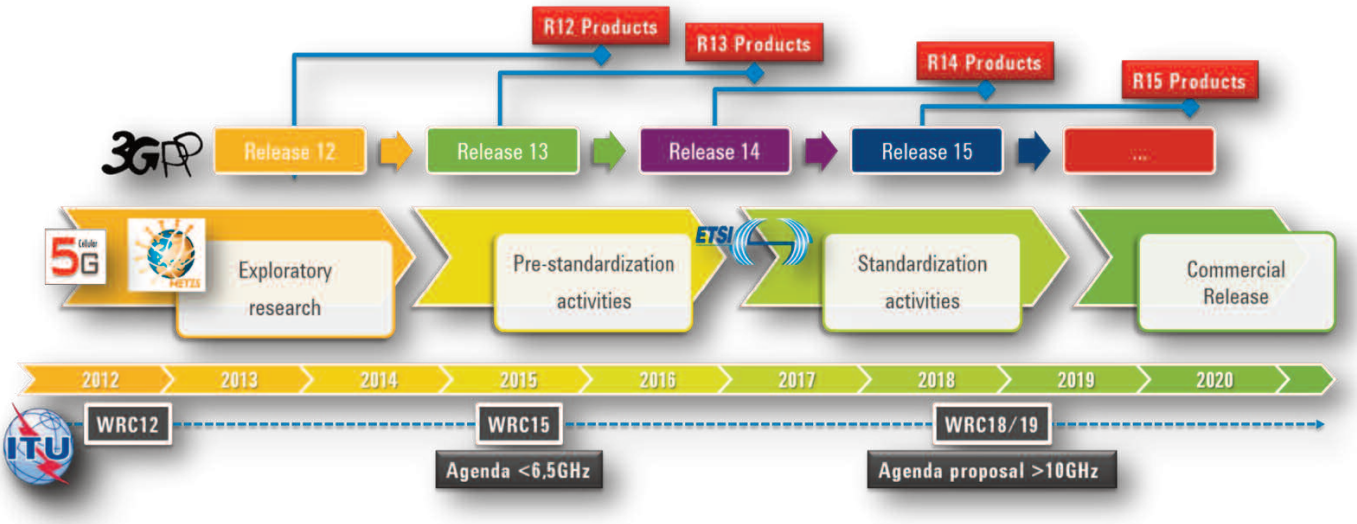

Getting the best out of DAS

Technology and techniques developed for dealing with network capacity variations need to work harmoniously to deliver real benefits. By John Spindler, Director of Product Management, TE Connectivity

Enhanced mobile infrastructure is often a requirement for large venues and outdoor areas. With large amounts of people all in one place using their mobile devices to send photos, access the internet or transmit video simultaneously, it becomes difficult for mobile operators to provide enough network capacity to facilitate these high demands.

Putting up one large antenna will not deliver enough performance in these situations. Rather, the coverage and capacity must be focused to each seating section. This means that the wireless solution designs needs to be heavily sectorised to deliver defined capacity to each group of seats and other service areas. Another requirement is that the design should minimise soft hand-off (SHO) from signals bleeding from one sector into another. As a result, the solution must be designed to accurately target the signal so as not to create overlap between sectors, or create isolation. Venue operators must also decide whether the coverage and capacity enhancement is intended to be permanent (as in a sports stadium) or temporary. With this in mind, operators are increasingly using distributed antenna systems (DAS) to divide facilities into sectors using discreet, remote antennas. However, like any investment, they are constantly looking at ways to get more efficiency out of capacity for DAS while conserving costs, to ensure users in areas of high density are able to take full advantage of their mobile devices.

There are three key trends at work that can lead to lower costs and higher service quality in mobile networks; base station hotels, fibre conservation technologies, and coordinated multipoint (CoMP) technology that improves service at the edges of the network. Whether used separately or together, these technologies will enable mobile operators to improve mobile services at large venues, to reduce churn, while maximising the use of their resources and keeping costs under control.

Three ways

With today's fibre-fed DAS, mobile operators can locate base stations and DAS head-ends miles away from the stadium or arena where the DAS is deployed. By combining base stations into a central 'hotel,' operators can avoid having to find the space required at the venues themselves. In addition, they can share the base stations' capacity among three different facilities, and reduce backhaul costs by backhauling traffic from one location instead of many. This model makes efficient use of base station resources while putting all of the base

stations in one location – instead of three – for easier access and maintenance.

The trend toward base station hotels feeding multiple large venues points to the need for lots of fibre to connect the base station hotel to a DAS. Typically, each DAS head end requires one to three fibre pairs, but there can be a dozen or more head ends in a base station hotel. Finding the fibre to transport the DAS traffic between the base station hotel and the DAS-serviced venues can be problematic.

One solution to this is a new type of 'muxponder,' which is a multiplexer and transponder in one unit that will take in three 3.072 gigabit per second (Gbit/s) feeds and multiplex them into a 9.8304Gbit/s transport over a single fibre pair. By combining three fibre pairs into one, the muxponder saves two-thirds of the fibre needed coming out of a base station hotel. These solutions are also perfect for neutral host architectures where it is necessary to transport full band, multi band RF to a designated service area, such as a stadium or an urban core, where there is high sectorisation and capacity strain on the network.

Base station hotels are becoming a popular strategy for mobile operators looking to get more efficiency out of base station capacity, while using DAS to serve large venues or urban cores. With muxponders, operators can extend these efficiencies out into the fibre network by slashing the amount of fibre needed to transport DAS traffic.

CoMP technology addresses service deficiencies for mobile devices at the edges of cells. When a mobile device is at the edge of the cell, the data rate drops off and the device starts looking to hand off its connection to the next cell. The signal levels are constantly changing and service is poor in these areas. The idea of CoMP is to get two or more cell sites to cooperate. One CoMP scenario is that the network can send the data to both base stations and, on a real-time basis, the base

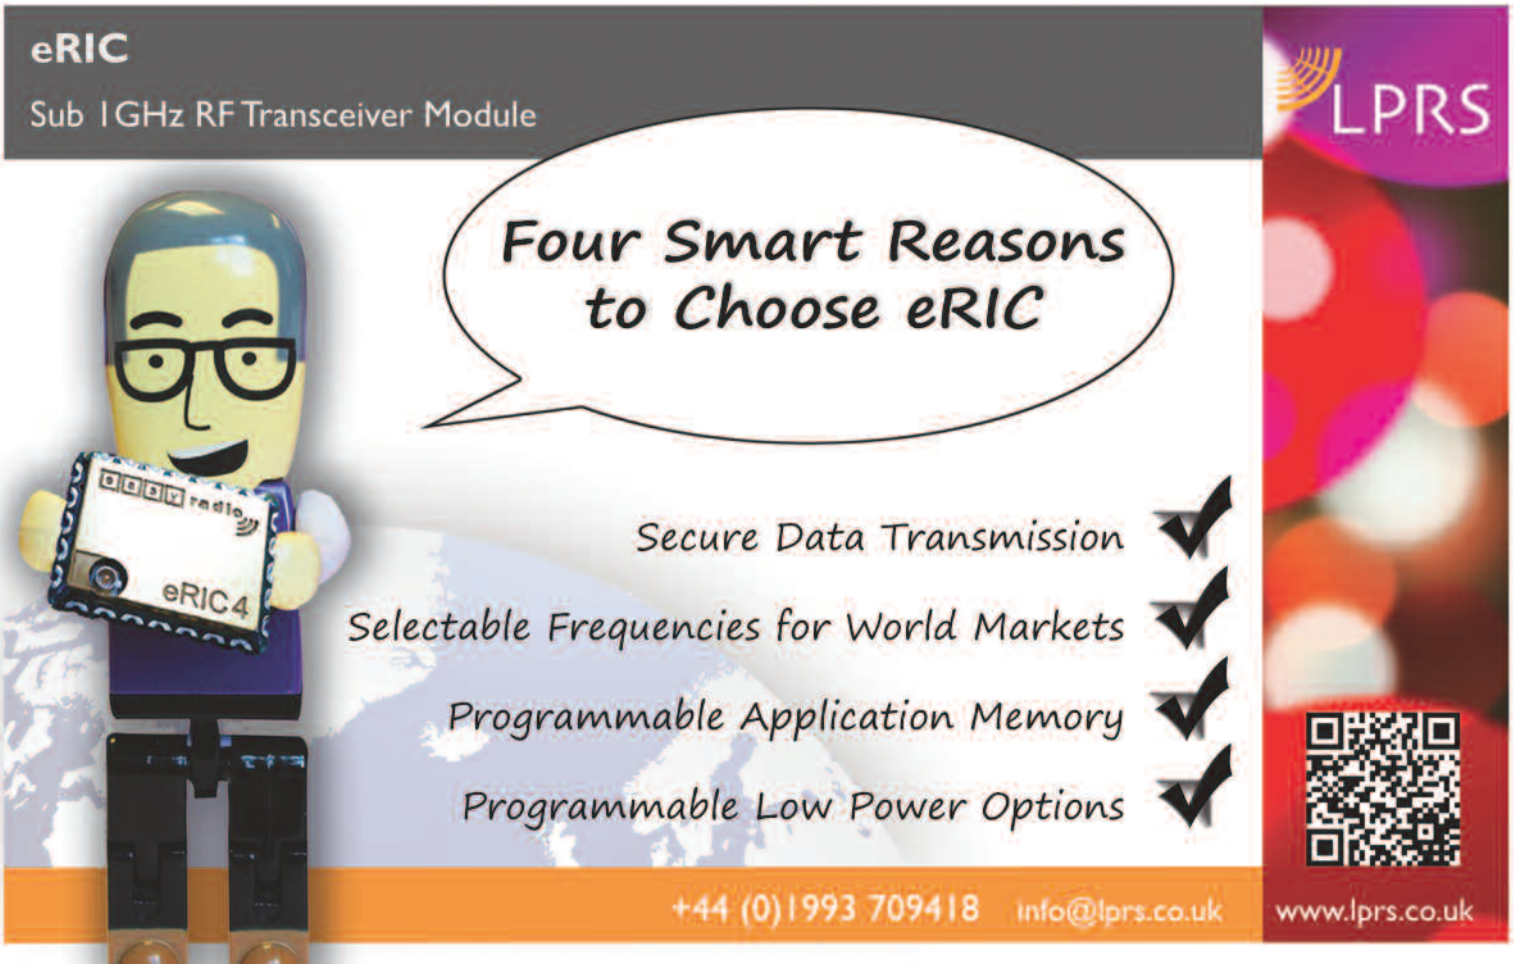

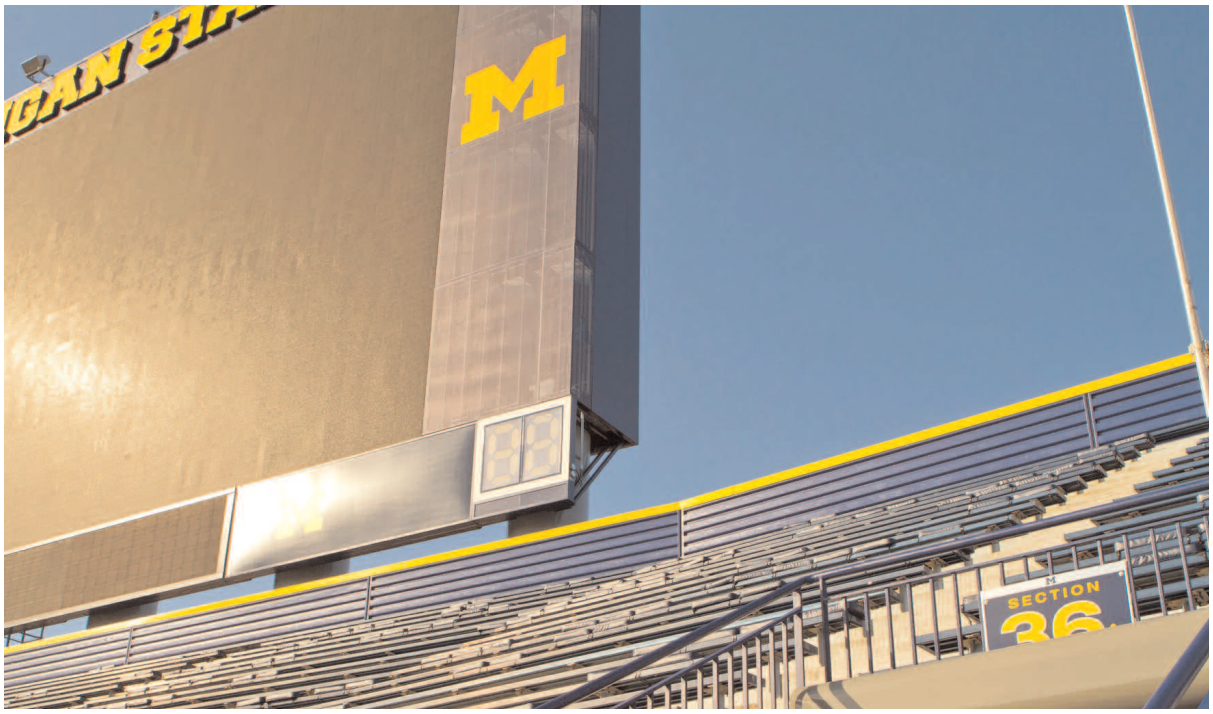

stations monitor the signal quality and can actually decide which of the base stations is in the best position to get that signal to the mobile device. In another form of CoMP, only one base station is transmitting but the other base stations in the area are aware of which time slots and frequency bands are being used to communicate, and they cooperate by saying 'we won't use that same block in our cell so we won't interfere with you'.

Cooperation

So, under CoMP, all of the base stations in an area identify which mobile devices are in the lowest signal areas and figure out how to cooperate to get the best throughput to that mobile either by reducing interference or by sharing the responsibility of transmitting. CoMP is included in the LTE portion of the 3GPP specifications beginning with Release 11. It is an optional enhancement to the LTE-Advanced (LTE-A) air interface technology that was introduced in 3GPP Release 10.

One of the main challenges with CoMP is backhaul. In order to coordinate their operations, cell sites must be connected to each other as well as to the core network, so this means that mobile operators will have to install a lot of expensive new backhaul equipment between cell sites. However, by pooling base stations in a hotel, the backhaul requirements are easily and cost-effectively managed.

With base stations housed in a central hotel, the base stations can coordinate their efforts easily and then transmit signals through a DAS to the desired area. In fact, mobile operators have already determined that the only way to make CoMP work in the real world is to pool base stations in a hotel and use remote radio heads or DAS to distribute their signals. With CoMP technology, cell sites can become smart and dynamic, improving service to devices at the edge of a cell and creating happier customers.

Ultimately, as the demand for mobile network services anytime, anywhere increases, mobile operators are under huge amounts of pressure to find ways to deliver better wireless coverage, particularly in densely populated venues or events. They want to improve services to reduce churn, but they also want to maximise the use of their resources and keep costs under control. Base station hotels, fibre conservation technologies and CoMP work separately or together to help operators achieve these goals and essentially make the most out of DAS.

Figure 1: High-level functional view of the UltraScale DSP48 slice

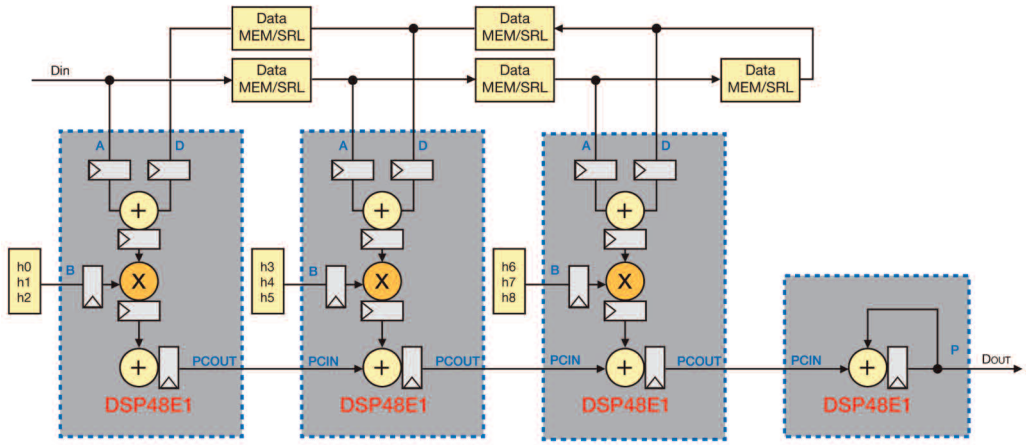

Figure 2: Implementation of a semi-parallel filter on 7 Series (a, above) and UltraScale (b, below) architectures

support for floating-point arithmetic. First, it is worth pointing out that the DSP48E2 can in effect support up to 28x18-bit or 27x19-bit signed multiplication, achieved by using the C input to process the additional bit.

This makes it possible to implement a 28x18-bit multiplier with a single DSP48E2 slice and 18 LUT/flip-flop pairs. The same applies for a 27x19-bit multiplier, using 27 additional LUT/flip-flop pairs. In both cases, convergent rounding of the result can still be supported through the W-mux.

A double-precision floating-point multiplication involves the integer product of the 53-bit unsigned mantissas of both operators. Although a 52-bit value (m) is stored in the double-precision floating-point representation, it describes the fractional part of the unsigned mantissa, and it is actually the normalised 1+m values, which need to be multiplied together, hence the additional bit required by the multiplication. Taking into account the fact that the MSBs of both 53-bit operands are equal to 1, and appropriately splitting the multiplication to optimally exploit the DSP48E2 26x17-bit unsigned multiplier and its improved capabilities (e.g., the true three-input 48-bit adder enabled by the W-mux), it can be shown that the 53x53-bit unsigned multiplication can be built with only six DSP48E2 slices and a minimal amount of external logic.

The 27x18 multiplier of the DSP48E2 is also very useful for applications based on fused data paths. The concept of a fused multiply-add operator has been recently added to the IEEE floating-point standard. Basically, it consists of building the floating-point operation A*B+C, without explicitly rounding, normalising and de-normalising the data between the multiplier and the adder. These functions are indeed very costly when using traditional floating-point arithmetic and account for the greatest part of the latency. This concept may be generalised to build sum-of-products operators, which are common in linear algebra (matrix product, Cholesky decomposition). Consequently, such an approach is quite efficient for applications where cost or latency are critical, while still requiring the accuracy and dynamic range of the floating-point representation. This is the case in radio DFE applications for which the digital pre-distortion functionality usually requires some hardware-acceleration support to improve the update rate of the nonlinear filter coefficients. You can then build one or more floating-point MAC engines in the FPGA fabric to assist the coefficient-estimation algorithm running in software (e.g. on one of the ARM Cortex-A9 cores of the Zynq SoC).

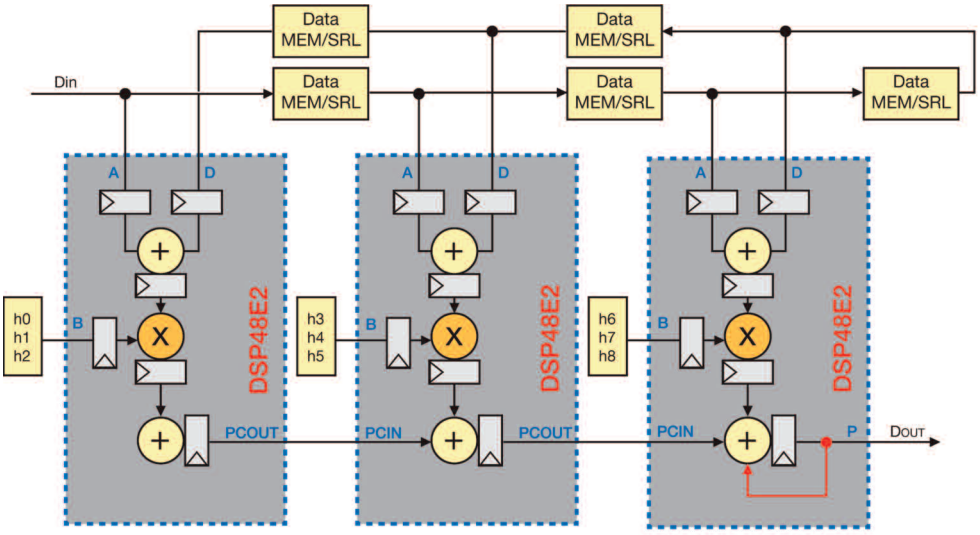

For such arithmetic structures, it has been shown that a slight increase of the mantissa width from 23 to 26 bits can provide even better accuracy compared with a true single-precision floating-point implementation, but with reduced latency and footprint. The UltraScale architecture is again well adapted for this purpose, since it takes only two DSP48 slices to build a single-precision fused multiplier, whereas three are required on 7 Series devices with additional fabric logic.

The pre-adder, integrated within the DSP48 slice in front of the multiplier, provides an efficient way to implement symmetric filters that are commonly used in DFE designs to realise the digital upconverter (DUC) and downconverter (DDC) functionality.

Fourth input

It is indisputably the addition of a fourth input operand to the ALU, through the extra W-mux multiplexer, which brings the most benefit for radio applications. This operand can typically save 10 to 20% of the DSP48 requirements for such designs compared with the same implementation on a 7 Series device.

The W-mux output can only be added within the ALU (subtraction is not permitted), and can be set dynamically as the content of either the C or P register or as a constant value, defined at FPGA configuration (e.g. the constant to be added for convergent or symmetric rounding of the DSP48 output), or simply forced to 0. This allows performing a true three-input operation when the multiplier is used, such as A*B+C+P, A*B+C+PCIN, A*B+P+PCIN, something that is not possible with the 7 Series architecture. Indeed, the multiplier stage generates the last two partial-product outputs, which are then added within the ALU to complete the operation. Therefore, when enabled, the multiplier uses two inputs of the ALU, and a three-input operation cannot be performed on 7 Series devices. Two of the most significant examples that benefit from this additional ALU input are semi-parallel filters and complex multiply-accumulate (MAC) operators.

Linear filters are the most common processing units of any DFE application. When integrating such functionality on Xilinx FPGAs, it is recommended, as far as possible, to implement multichannel filters for which the composite sampling rate (defined as the product of the number of channels by the common signal-sampling frequency of each channel) is equal to the clock rate at which the design is running. In a so-called parallel architecture, each DSP48 slice supports a single filter coefficient per data channel, which greatly simplifies the control logic and hence minimises the design resource utilisation.

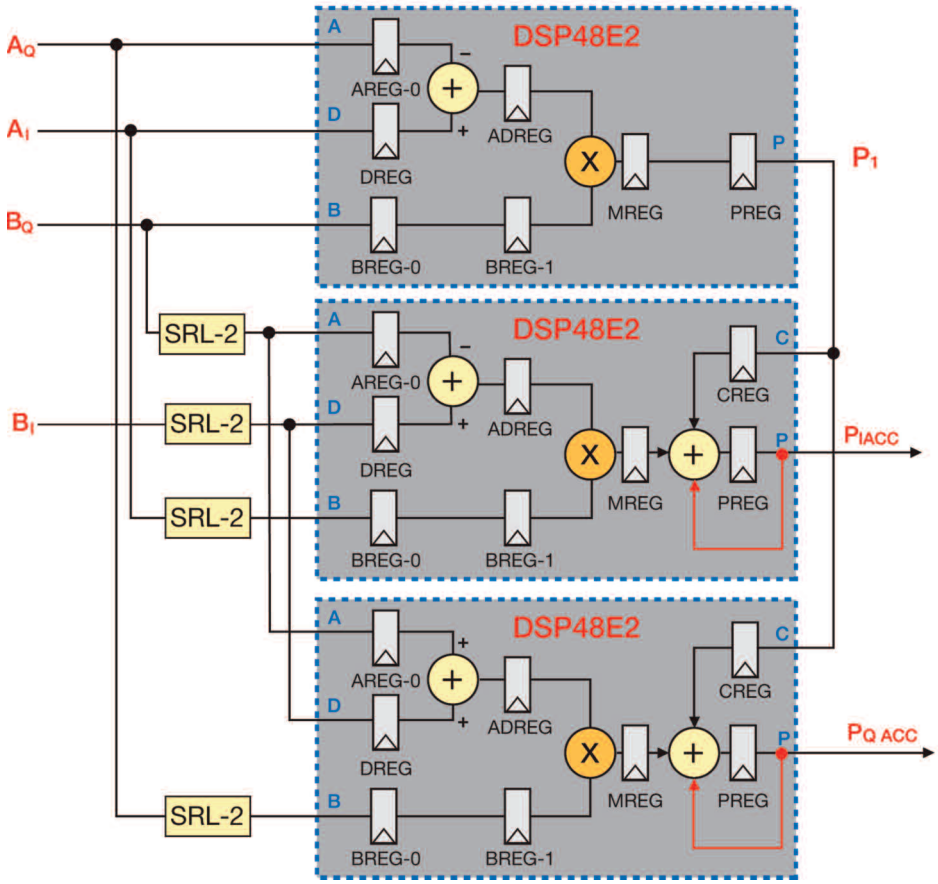

Figure 3: Implementation of a complex MAC on 7 Series and UltraScale architectures

However, with higher clock-rate capabilities (for example, more than 500MHz on lowest-

speed-grade UltraScale devices), and for filters running at a relatively low sampling rate, it is often the case that the clock rate can be selected as a multiple of the composite sampling rate. It's desirable to increase the clock rate as much as possible to further reduce the design footprint, as well as the power consumption. In such situations, a semi-parallel architecture is built where each DSP48 processes K coefficients per channel, where K is the ratio between the clock rate and the composite sampling rate. The most efficient implementation then consists of splitting the filter into its K phases, each DSP48 processing a specific coefficient of these K phases.

At each clock cycle, the successive phases of the filter output are computed and need to be accumulated together to form an output sample (once every K cycle). Consequently, an additional accumulator is required at the filter output compared with a parallel implementation. This full-precision accumulator works on a large data width, equal to bS+bC+bF, where bS and bC are respectively the bit widths of the data samples and coefficients, and bF=Log2N is the filter bit growth, N being the total number of coefficients. Normal practice is therefore to implement the accumulator within a DSP48 slice to ensure support for the highest clock rate while minimising footprint and power.

It should be noted that semi-parallel architectures can be derived for any type of filter: single-rate, integer or fractional-rate interpolation and decimation. Figure 2 shows a simplified block diagram for both 7 Series and UltraScale implementations. It clearly highlights the advantage of the UltraScale solution, since the phase accumulator is absorbed by the last DSP48 slice thanks to the W-mux capability.

It is well known that you can rewrite the equation of a complex product to use only three real multiplications. By exploiting the built-in pre-adder, you can implement a complex multiplier with three DSP48s only; one to compute P1 and the other two to handle the PI and PQ outputs. Depending on the latency requirements, which also dictate the speed performance, some logic needs to be added to balance the delays between the different data paths. To get maximal speed support, the DSP48 must be fully pipelined, which results in an overall latency of six cycles for the operator. A two-cycle delay line is consequently added on each input to correctly align the real and imaginary data paths. Those are implemented with four SRL2 per input bit, which are in effect packed into two LUTs by taking advantage of the SRL compression capabilities.

The complex MAC is finally completed by adding an accumulator on each of the PI and PQ outputs. Again this accumulator works on large data widths and is therefore better integrated within a DSP48 slice. The corresponding implementation for an UltraScale device is shown in Figure 3, which demonstrates the benefit of the W-mux integration. The PI and PQ DSP48E2 slices absorb the accumulators, with 40% resource savings. It is worth mentioning that the latency is also reduced, which may be beneficial for some applications.

Using a similar construction, you can build a complex filter (one with complex data and coefficients) with three real filters. The real and imaginary parts of the input signal are fed into two real filters, with coefficients derived respectively as the difference and sum of the imaginary and real parts of the filter coefficients. The third filter processes the sum of the input real and imaginary parts in parallel, using the real part of the coefficients.

The outputs of these three filters are finally combined to generate the real and imaginary components of the output, which can again benefit from the W-mux, when parallel filters need to be built, which is typically the case for the equalisers used in DFE applications.

Bringing the real world into the lab

David Hill, VP EMEA at Spirent, explains how new approaches to end-to-end testing for wireless communications can save time and money by modelling realistic test conditions even in the development stages

Wireless technology is, by default, complicated and every several years, a surge of new technologies presents a dramatic increase in complexity. Today, major innovations in the mobile device, Radio Access Network (RAN), core network and services are all hitting at once, driving wireless technology to unprecedented levels of complexity and presenting significant challenges in delivering the expected Quality of Service (QoS) to every mobile subscriber.

Wireless communication has never been a simple matter, but every few years a number of technical innovations come together and the level of complexity takes a quantum leap. This is happening right now. Mobile devices come in an awesome range of variations based on a range of chipsets, radios, displays, batteries and operating systems. As well as these innovations in the device, new developments in the Radio Access Network (RAN), core network and service layer all impact wireless technology and raise it to unprecedented levels of complexity (Figure 1).

With increased complexity comes the risk that the test procedures of the past no longer identify issues between today□s more complex interactions and new technology deployments. This means that testing and quality validation efforts need to be more flexible and faster to ensure the issues are identified before services and equipment go live. Without faster and more precise testing, it is impossible to be sure that new technology, devices and infrastructure will deliver services that meet end-user expectations in the live network.

This pressure to improve testing and validation has driven leading industry players to rethink the whole process of bringing new services, devices and infrastructure to market. One key priority is to isolate problems as early as possible. As we proceed further down the development path, more work is entailed in correcting design errors, and as the process moves from a controlled lab environment to a dynamic live network, even more time and resources are needed to isolate and correct problems. When testing and validation move from the lab to the live network, cost and delays caused by emerging problems rise dramatically. Identifying problems earlier will save much time and money.

New levels of testing realism

The latest test solutions, designed specifically to address these new demands, enable the live network to be modelled with unprecedented realism and accuracy in a controlled lab environment. By accurately and realistically emulating real-world conditions in the lab, many problems that only appear in live network testing can be successfully identified and resolved during lab tests; slashing the resources and time needed to get new devices, services and infrastructure launched.

These solutions increasingly use live network

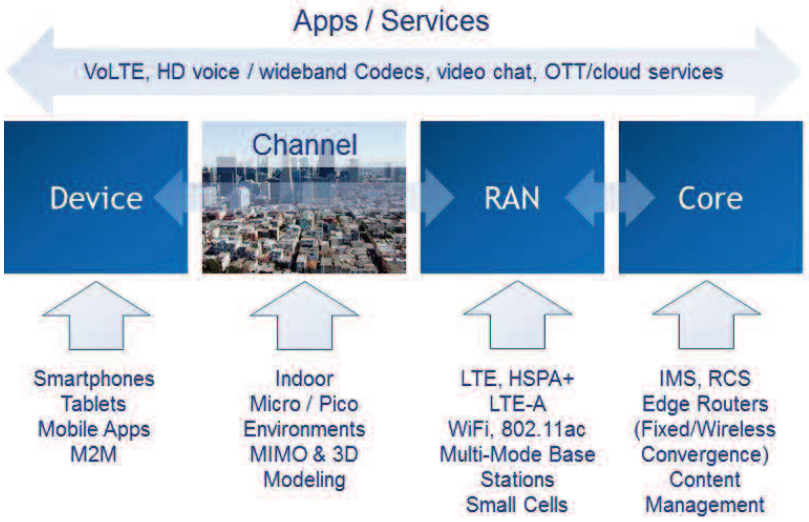

Figure 1: A wave of innovations in wireless technology is leading to unprecedented complexity.

measurements of various sources from protocol captures to RF signal data to drive sophisticated emulation engines. The current generation solutions use actual live network drive tests to enable their channel emulation engines to accurately model the dynamic RF conditions of multi-cell live networks.

In a live network, RF transmissions are incredibly dynamic, consisting of fast and slow fading signals from hundreds of cells. In addition to all those RF signals, there is noise and extra-system interference. Furthermore, both HSPA and LTE rely on multiple-input multiple-output (MIMO) antenna transmission to improve system performance.

Modelling live wireless environments like these in the lab raises several key challenges. First, the cost and complexity of base stations limits the number of cells which can be used for lab tests; typically from 1-4 cells will be available for lab tests compared with 10-100 cells along a test route in the live network. So the signal variations measured across a large number of cells in the live network need to be mapped onto a much smaller number of lab-based cells, while still capturing the essential variability of the desired and interfering signals.

A second key challenge is to accurately model MIMO configurations in the lab, which may not yet exist in the live-network. For example, an LTE live network may only have 2x2 MIMO deployed, whereas lab tests may require testing for 4x4 or other future MIMO configurations. Furthermore, drive test measurement equipment may not capture all aspects of the wireless channel with sufficient resolution to enable an accurate channel model to be developed. To solve these challenges, lab-based channel models must begin with as much data as possible from live network drive tests and then model the additional channel aspects that aren't available in the drive test.

Thirdly, the channel models developed from the drive test data must be converted into an import format that a channel emulator can use to create a real-time wireless channel that mimics live network conditions. This import format must take into account the specifics of the experimental configuration, including the number of RF antennas used by the base station and mobile device, the bi-directional or unidirectional nature of the test and other factors.

A radical approach is needed to address all of these challenges. Spirent's solution is based on a "Virtual Drive Test Conversion Tool" (VDT-CT) that converts live network drive test data from multiple commercial field capture tools into channel models which may be imported into our channel emulators.

VDT-CT addresses the challenge of modelling live wireless environments in the lab in three main stages: Filtering live network drive tests; Mapping live network data to lab-based cells; and, Generating a channel emulator configuration file. Each of these is based on algorithms that are technology independent, applying equally to CDMA, UMTS & LTE technology families.

The main purpose of the filtering algorithm is to remove fast-fading effects from the live network drive test data. This leaves a processed set of signals from each base station which includes only slow fading effects, i.e. signal variations due to shadowing of buildings or other obstructions between the base station and the mobile device.

Using only these slow fading effects from the drive test allows the channel emulator to apply any supported MIMO configuration on top of the data. The channel emulator uses the slow-fading signal as a baseline for all MIMO transmission paths, then adds fast-fading and Doppler effects with a user-defined level of signal correlation on top of the slow-fading signal.

Fast-fading effects are filtered using the well-known and widely used wavelength-based averaging approach. Signals are averaged over a 40□ distance (at 2GHz, 40□=6m). VDT-CT also incorporates several algorithms for mapping signals from live network cells to lab-based cells; each has unique characteristics suited to specific testing needs. More information on the specifics of each algorithm can be obtained from the Spirent VDT-CT user manual. Figure 2 is a sample of the Preservation (Normal) algorithm in action, showing a typical UMTS live network drive test, during which more than 76 cells were observed.

The last step in the live network modelling process is to take the processed drive test data and convert it into a form that the channel emulator can use. In order to do this, VDT-CT needs to know key aspects of the lab-based test configuration such as: Technology (e.g. LTE FDD, LTE TDD); Band, and; Fader connection type (Uplink and downlink, RF channel mapping, MIMO channel mapping and order (e.g. 2x2), Bi-directional or unidirectional). Using this information, VDT-CT creates a live network channel model file that can be imported into the channel emulator and used to emulate the live network environment in the lab.

Building in realism

It is a clear advantage to be able to emulate real world test conditions in the laboratory throughout the design process, and it is equally important to deliver these tests in a simple, repeatable and easy-to-use manner.

Spirent does this by providing automation solutions for end-to-end QoE measurements using a mix of real off-the-shelf devices and, optionally, real base station or core network infrastructure if needed. This automated virtual drive test can cover device testing, RAN validation and pre-launch testing of new services. We also offer a specific mobile device test solution for automating carrier acceptance tests for mobile devices.

The ultimate test of any wireless network must lie in the hands of the user, but we are faced with such a range of devices and performance across use cases and applications that the challenge of ensuring an all-around enjoyable mobile experience is growing. Launching wireless services that delight users demands the development and execution of comprehensive and robust test scenarios well in advance of their launch. By the time they launch, the cost of correcting mistakes may be prohibitive both in terms of fault rectification and damage to reputation.

The solution is to emulate as much of that final test as simply and realistically as possible early in the development process. The latest generation of test solutions not only achieve this in the laboratory with new levels of precision, they also deliver these highly sophisticated tests in a format that is easy to use, automated and highly flexible to allow nimble test adaptation to fast-moving changes in technology and network deployment.

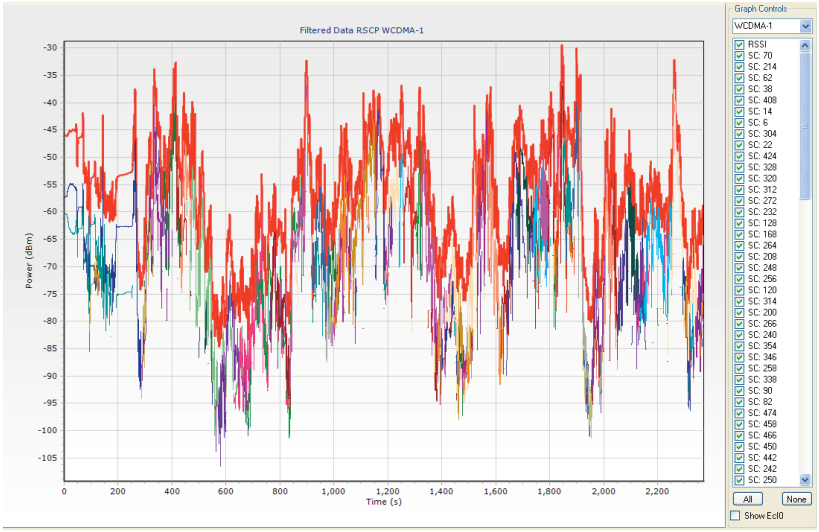

Figure 2: RSCP data from a drive test of a live UMTS network. 76 cells were observed during the test.

PCIM
SOUTH AMERICA

International Conference and Exhibition
for Power Electronics, Intelligent Motion,
Renewable Energy and Energy Management

São Paulo, 14 – 15 October 2014

Power On!

pcim-southamerica.com

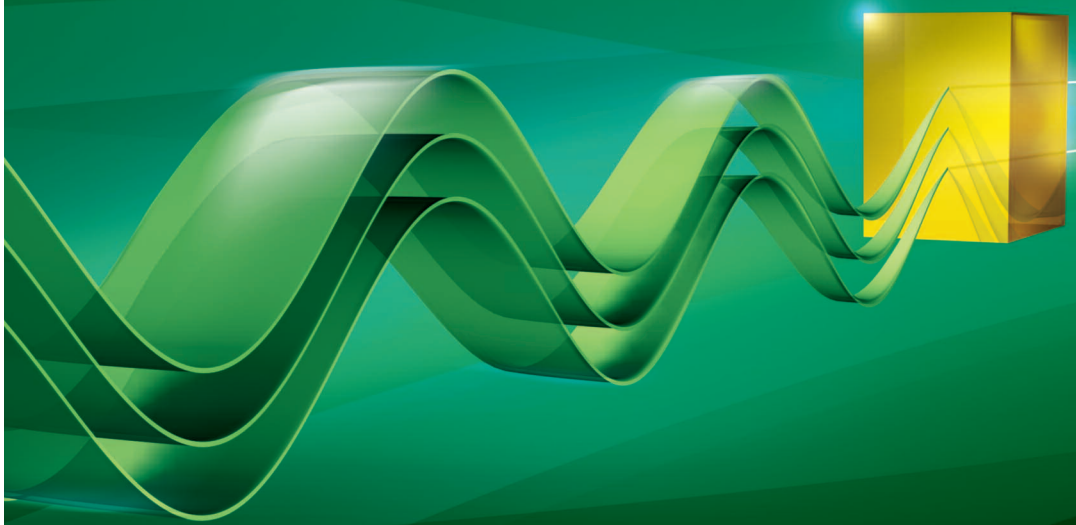

Mesago
PCIM